# Workload APIs

CronJob [batch/v1]

DaemonSet [apps/v1]

Deploymen

Job [batch/v1]

Pod [v1]

Replicaset |

ReplicationController [v1]

StatefulSet [apps/v1]

# CronJob [batch/v1]

**Description**

CronJob represents the configuration of a single cron job.

**Type**

`object`

## Specification

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources |
| `kind` | `string` | Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | CronJobSpec describes how the job execution will look like and when it will actually run. |
| `status` | `object` | CronJobStatus represents the current state of a cron job. |

### .spec

**Description**

CronJobSpec describes how the job execution will look like and when it will actually run.

**Type**

`object`

**Required**

`schedule` `jobTemplate`

| Property | Type | Description |
| --- | --- | --- |
| `concurrencyPolicy` | `string` | Specifies how to treat concurrent executions of a Job. Valid values are: |

| Property | Type | Description |
| --- | --- | --- |
| | | • "Allow" (default): allows CronJobs to run concurrently; - "Forbid": forbids concurrent runs, skipping next run if previous run hasn't finished yet; - "Replace": cancels currently running job and replaces it with a new one<br><br>Possible enum values:<br><br>• `"Allow"` allows CronJobs to run concurrently.<br>• `"Forbid"` forbids concurrent runs, skipping next run if previous hasn't finished yet.<br>• `"Replace"` cancels currently running job and replaces it with a new one. |
| `failedJobsHistoryLimit` | `integer` | The number of failed finished jobs to retain. Value must be non-negative integer. Defaults to 1. |
| `jobTemplate` | `object` | JobTemplateSpec describes the data a Job should have when created from a template |
| `schedule` | `string` | The schedule in Cron format, see https://en.wikipedia.org/wiki/Cron ↗. |
| `startingDeadlineSeconds` | `integer` | Optional deadline in seconds for starting the job if it misses scheduled time for any reason. Missed jobs executions will be counted as failed ones. |
| `successfulJobsHistoryLimit` | `integer` | The number of successful finished jobs to retain. Value must be non-negative integer. Defaults to 3. |
| `suspend` | `boolean` | This flag tells the controller to suspend subsequent executions, it does not apply to already started executions. Defaults to false. |
| `timeZone` | `string` | The time zone name for the given schedule, see https://en.wikipedia.org/wiki/List_of_tz_database_time_zones ↗. If not specified, this will default to the time zone of the kube-controller-manager process. The set of valid time zone names and the time zone offset is loaded from the system-wide time zone database by the API server during CronJob validation and the controller manager during execution. If no system-wide time zone database can be found a bundled version of the database is used instead. If the time zone name becomes invalid during the lifetime of a CronJob or due to a change in host configuration, the controller will stop creating new new Jobs and will create a system event with the reason UnknownTimeZone. More information can be found in https://kubernetes.io/docs/concepts/workloads/controllers/cron-jobs/#time-zones ↗ |

## .spec.jobTemplate

### Description

JobTemplateSpec describes the data a Job should have when created from a template

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | JobSpec describes how the job execution will look like. |

## .spec.jobTemplate.spec

### Description

JobSpec describes how the job execution will look like.

### Type

`object`

### Required

`template`

| Property | Type | Description |
| --- | --- | --- |
| `activeDeadlineSeconds` | `integer` | Specifies the duration in seconds relative to the startTime that the job may be continuously active before the system tries to terminate it; value must be positive integer. If a Job is suspended (at creation or through an update), this timer will effectively be stopped and reset when the Job is resumed again. |
| `backoffLimit` | `integer` | Specifies the number of retries before marking this job failed. Defaults to 6 |
| `backoffLimitPerIndex` | `integer` | Specifies the limit for the number of retries within an index before marking this index as failed. When enabled the number of failures per index is kept in the pod's batch.kubernetes.io/job-index-failure-count annotation. It can only be set when Job's completionMode=Indexed, and the Pod's restart policy is Never. The field is immutable. This field is beta-level. It can be used when the `JobBackoffLimitPerIndex` feature gate is enabled (enabled by default). |
| `completionMode` | `string` | completionMode specifies how Pod completions are tracked. It can be `NonIndexed` (default) or `Indexed`.<br><br>`NonIndexed` means that the Job is considered complete when there have been .spec.completions successfully completed Pods. Each Pod completion is homologous to each other.<br><br>`Indexed` means that the Pods of a Job get an associated completion index from 0 to (.spec.completions - 1), available in the annotation batch.kubernetes.io/job-completion-index. The Job is considered complete when there is one successfully completed Pod for each index. When value is `Indexed`, .spec.completions must be specified and `.spec.parallelism` must be less than or equal to 10^5. In addition, The Pod name takes the form `$(job-name)-$(index)-$(random-string)`, the Pod hostname takes the form `$(job-name)-$(index)`.<br><br>More completion modes can be added in the future. If the Job controller observes a mode that it doesn't recognize, which is possible during upgrades due to version skew, the controller skips updates for the Job. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"Indexed"` is a Job completion mode. In this mode, the Pods of a Job get an associated completion index from 0 to (.spec.completions - 1). The Job is considered complete when a Pod completes for each completion index.<br>- `"NonIndexed"` is a Job completion mode. In this mode, the Job is considered complete when there have been .spec.completions successfully completed Pods. Pod completions are homologous to each other. |
| `completions` | `integer` | Specifies the desired number of successfully finished pods the job should be run with. Setting to null means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/ ↗ |
| `managedBy` | `string` | ManagedBy field indicates the controller that manages a Job. The k8s Job controller reconciles jobs which don't have this field at all or the field value is the reserved string `kubernetes.io/job-controller`, but skips reconciling Jobs with a custom value for this field. The value must be a valid domain-prefixed path (e.g. acme.io/foo) - all characters before the first "/" must be a valid subdomain as defined by RFC 1123. All characters trailing the first "/" must be valid HTTP Path characters as defined by RFC 3986. The value cannot exceed 63 characters. This field is immutable.<br><br>This field is beta-level. The job controller accepts setting the field when the feature gate JobManagedBy is enabled (enabled by default). |
| `manualSelector` | `boolean` | manualSelector controls generation of pod labels and pod selectors. Leave `manualSelector` unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see `manualSelector=true` in jobs that were created with the old `extensions/v1beta1` API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector ↗ |
| `maxFailedIndexes` | `integer` | Specifies the maximal number of failed indexes before marking the Job as failed, when backoffLimitPerIndex is set. Once the number of failed indexes exceeds this number the entire Job is marked as Failed and its execution is terminated. When left as null the job continues execution of all of its indexes and is marked with the `Complete` Job condition. It can only be specified when backoffLimitPerIndex is set. It can be null or up to completions. It is required and must be less than or equal to 10^4 when is completions greater than 10^5. This field is beta-level. It can be used when the `JobBackoffLimitPerIndex` feature gate is enabled (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `parallelism` | `integer` | Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when ((.spec.completions - .status.successful) < .spec.parallelism), i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/ ↗ |
| `podFailurePolicy` | `object` | PodFailurePolicy describes how failed pods influence the backoffLimit. |
| `podReplacementPolicy` | `string` | podReplacementPolicy specifies when to create replacement Pods. Possible values are: - TerminatingOrFailed means that we recreate pods when they are terminating (has a metadata.deletionTimestamp) or failed.<br>• Failed means to wait until a previously created Pod is fully terminated (has phase Failed or Succeeded) before creating a replacement Pod.<br>When using podFailurePolicy, Failed is the the only allowed value. TerminatingOrFailed and Failed are allowed values when podFailurePolicy is not in use. This is an beta field. To use this, enable the JobPodReplacementPolicy feature toggle. This is on by default.<br>Possible enum values:<br>• `"Failed"` means to wait until a previously created Pod is fully terminated (has phase Failed or Succeeded) before creating a replacement Pod.<br>• `"TerminatingOrFailed"` means that we recreate pods when they are terminating (has a metadata.deletionTimestamp) or failed. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `successPolicy` | `object` | SuccessPolicy describes when a Job can be declared as succeeded based on the success of some indexes. |
| `suspend` | `boolean` | suspend specifies whether the Job controller should create Pods or not. If a Job is created with suspend set to true, no Pods are created by the Job controller. If a Job is suspended after creation (i.e. the flag goes from false to true), the Job controller will delete all active Pods associated with this Job. Users must design their workload to gracefully handle this. Suspending a Job will reset the StartTime field of the Job, effectively resetting the ActiveDeadlineSeconds timer too. Defaults to false. |
| `template` | `object` | PodTemplateSpec describes the data a pod should have when created from a template |
| `ttlSecondsAfterFinished` | `integer` | ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this |

| Property | Type | Description |
| --- | --- | --- |
| | | field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. |

## .spec.jobTemplate.spec.podFailurePolicy

**Description**

PodFailurePolicy describes how failed pods influence the backoffLimit.

**Type**

`object`

**Required**

`rules`

| Property | Type | Description |
| --- | --- | --- |
| `rules` | `array` | A list of pod failure policy rules. The rules are evaluated in order. Once a rule matches a Pod failure, the remaining of the rules are ignored. When no rule matches the Pod failure, the default handling applies - the counter of pod failures is incremented and it is checked against the backoffLimit. At most 20 elements are allowed. |

## .spec.jobTemplate.spec.podFailurePolicy.rules

**Description**

A list of pod failure policy rules. The rules are evaluated in order. Once a rule matches a Pod failure, the remaining of the rules are ignored. When no rule matches the Pod failure, the default handling applies - the counter of pod failures is incremented and it is checked against the backoffLimit. At most 20 elements are allowed.

**Type**

`array`

## .spec.jobTemplate.spec.podFailurePolicy.rules[]

**Description**

PodFailurePolicyRule describes how a pod failure is handled when the requirements are met. One of onExitCodes and onPodConditions, but not both, can be used in each rule.

**Type**

`object`

**Required**

`action`

| Property | Type | Description |
| --- | --- | --- |
| `action` | `string` | Specifies the action taken on a pod failure when the requirements are satisfied. Possible values are:<br>• FailJob: indicates that the pod's job is marked as Failed and all running pods are terminated.<br>• FailIndex: indicates that the pod's index is marked as Failed and will not be restarted. This value is beta-level. It can be used when the `JobBackoffLimitPerIndex` feature gate is enabled (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| | | • Ignore: indicates that the counter towards the .backoffLimit is not incremented and a replacement pod is created.<br>• Count: indicates that the pod is handled in the default way - the counter towards the .backoffLimit is incremented. Additional values are considered to be added in the future. Clients should react to an unknown action by skipping the rule.<br><br>Possible enum values:<br><br>• `"Count"` This is an action which might be taken on a pod failure - the pod failure is handled in the default way - the counter towards .backoffLimit, represented by the job's .status.failed field, is incremented.<br>• `"FailIndex"` This is an action which might be taken on a pod failure - mark the Job's index as failed to avoid restarts within this index. This action can only be used when backoffLimitPerIndex is set. This value is beta-level.<br>• `"FailJob"` This is an action which might be taken on a pod failure - mark the pod's job as Failed and terminate all running pods.<br>• `"Ignore"` This is an action which might be taken on a pod failure - the counter towards .backoffLimit, represented by the job's .status.failed field, is not incremented and a replacement pod is created. |
| `onExitCodes` | `object` | PodFailurePolicyOnExitCodesRequirement describes the requirement for handling a failed pod based on its container exit codes. In particular, it lookups the .state.terminated.exitCode for each app container and init container status, represented by the .status.containerStatuses and .status.initContainerStatuses fields in the Pod status, respectively. Containers completed with success (exit code 0) are excluded from the requirement check. |
| `onPodConditions` | `array` | Represents the requirement on the pod conditions. The requirement is represented as a list of pod condition patterns. The requirement is satisfied if at least one pattern matches an actual pod condition. At most 20 elements are allowed. |

## .spec.jobTemplate.spec.podFailurePolicy.rules[].onExitCodes

**Description**

PodFailurePolicyOnExitCodesRequirement describes the requirement for handling a failed pod based on its container exit codes. In particular, it lookups the .state.terminated.exitCode for each app container and init container status, represented by the .status.containerStatuses and .status.initContainerStatuses fields in the Pod status, respectively. Containers completed with success (exit code 0) are excluded from the requirement check.

**Type**

`object`

**Required**

`operator` `values`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Restricts the check for exit codes to the container with the specified name. When null, the rule applies to all containers. When specified, it should match one the container or initContainer names in the pod template. |
| `operator` | `string` | Represents the relationship between the container exit code(s) and the specified values. Containers completed with success (exit code 0) are excluded from the requirement check. Possible values are: |

| Property | Type | Description |
| --- | --- | --- |
| | | • In: the requirement is satisfied if at least one container exit code (might be multiple if there are multiple containers not restricted by the 'containerName' field) is in the set of specified values.<br>• NotIn: the requirement is satisfied if at least one container exit code (might be multiple if there are multiple containers not restricted by the 'containerName' field) is not in the set of specified values. Additional values are considered to be added in the future. Clients should react to an unknown operator by assuming the requirement is not satisfied.<br><br>Possible enum values:<br>• `"In"`<br>• `"NotIn"` |
| `values` | `array` | Specifies the set of values. Each returned container exit code (might be multiple in case of multiple containers) is checked against this set of values with respect to the operator. The list of values must be ordered and must not contain duplicates. Value '0' cannot be used for the In operator. At least one element is required. At most 255 elements are allowed. |

## .spec.jobTemplate.spec.podFailurePolicy.rules[].onExitCodes.values

**Description**

Specifies the set of values. Each returned container exit code (might be multiple in case of multiple containers) is checked against this set of values with respect to the operator. The list of values must be ordered and must not contain duplicates. Value '0' cannot be used for the In operator. At least one element is required. At most 255 elements are allowed.

**Type**

`array`

## .spec.jobTemplate.spec.podFailurePolicy.rules[].onExitCodes.values[]

**Type**

`integer`

## .spec.jobTemplate.spec.podFailurePolicy.rules[].onPodConditions

**Description**

Represents the requirement on the pod conditions. The requirement is represented as a list of pod condition patterns. The requirement is satisfied if at least one pattern matches an actual pod condition. At most 20 elements are allowed.

**Type**

`array`

## .spec.jobTemplate.spec.podFailurePolicy.rules[].onPodConditions[]

**Description**

PodFailurePolicyOnPodConditionsPattern describes a pattern for matching an actual pod condition type.

**Type**

`object`

**Required**

`type` `status`

| Property | Type | Description |
| --- | --- | --- |
| `status` | `string` | Specifies the required Pod condition status. To match a pod condition it is required that the specified status equals the pod condition status. Defaults to True. |
| `type` | `string` | Specifies the required Pod condition type. To match a pod condition it is required that specified type equals the pod condition type. |

## .spec.jobTemplate.spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.successPolicy

**Description**

SuccessPolicy describes when a Job can be declared as succeeded based on the success of some indexes.

**Type**

`object`

**Required**

`rules`

| Property | Type | Description |
| --- | --- | --- |
| `rules` | `array` | rules represents the list of alternative rules for the declaring the Jobs as successful before `.status.succeeded >= .spec.completions`. Once any of the rules are met, the "SucceededCriteriaMet" condition is added, and the lingering pods are removed. The terminal state for such a Job has the "Complete" condition. Additionally, these rules are evaluated in order; Once the Job meets one of the rules, other rules are ignored. At most 20 elements are allowed. |

## .spec.jobTemplate.spec.successPolicy.rules

**Description**

rules represents the list of alternative rules for the declaring the Jobs as successful before `.status.succeeded >= .spec.completions`. Once any of the rules are met, the "SucceededCriteriaMet" condition is added, and the lingering pods are removed. The terminal state for such a Job has the "Complete" condition. Additionally, these rules are evaluated in order; Once the Job meets one of the rules, other rules are ignored. At most 20 elements are allowed.

**Type**

`array`

## .spec.jobTemplate.spec.successPolicy.rules[]

**Description**

SuccessPolicyRule describes rule for declaring a Job as succeeded. Each rule must have at least one of the "succeededIndexes" or "succeededCount" specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `succeededCount` | `integer` | succeededCount specifies the minimal required size of the actual set of the succeeded indexes for the Job. When succeededCount is used along with succeededIndexes, the check is constrained only to the set of indexes specified by succeededIndexes. For example, given that succeededIndexes is "1-4", succeededCount is "3", and completed indexes are "1", "3", and "5", the Job isn't declared as succeeded because only "1" and "3" indexes are considered in that rules. When this field is null, this doesn't default to any value and is never evaluated at any time. When specified it needs to be a positive integer. |
| `succeededIndexes` | `string` | succeededIndexes specifies the set of indexes which need to be contained in the actual set of the succeeded indexes for the Job. The list of indexes must be within 0 to ".spec.completions-1" and must not contain duplicates. At least one element is required. The indexes are represented as intervals separated by commas. The intervals can be a decimal integer or a pair of decimal integers separated by a hyphen. The number are listed in represented by the first and last element of the series, separated by a hyphen. For example, if the completed indexes are 1, 3, 4, 5 and 7, they are represented as "1,3-5,7". When this field is null, this field doesn't default to any value and is never evaluated at any time. |

## .spec.jobTemplate.spec.template

**Description**

PodTemplateSpec describes the data a pod should have when created from a template

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PodSpec is a description of a pod. |

## .spec.jobTemplate.spec.template.spec

### Description

PodSpec is a description of a pod.

### Type

`object`

### Required

`containers`

| Property | Type | Description |
| --- | --- | --- |
| `activeDeadlineSeconds` | `integer` | Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer. |
| `affinity` | `object` | Affinity is a group of affinity scheduling rules. |
| `automountServiceAccountToken` | `boolean` | AutomountServiceAccountToken indicates whether a service account token should be automatically mounted. |
| `containers` | `array` | List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated. |
| `dnsConfig` | `object` | PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy. |
| `dnsPolicy` | `string` | Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.<br><br>Possible enum values:<br>- `"ClusterFirst"` indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"ClusterFirstWithHostNet"` indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>• `"Default"` indicates that the pod should use the default (as determined by kubelet) DNS settings.<br>• `"None"` indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig. |
| `enableServiceLinks` | `boolean` | EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching traditional container linking syntax. Optional: Defaults to true. |
| `ephemeralContainers` | `array` | List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. |
| `hostAliases` | `array` | HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. |
| `hostIPC` | `boolean` | Use the host's ipc namespace. Optional: Default to false. |
| `hostNetwork` | `boolean` | Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false. |
| `hostPID` | `boolean` | Use the host's pid namespace. Optional: Default to false. |
| `hostUsers` | `boolean` | Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new userns is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature. |
| `hostname` | `string` | Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value. |
| `imagePullSecrets` | `array` | ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `initContainers` | `array` | List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/ |
| `nodeName` | `string` | NodeName indicates in which node this pod is scheduled. If empty, this pod is a candidate for scheduling by the scheduler defined in schedulerName. Once this field is set, the kubelet for this node becomes responsible for the lifecycle of this pod. This field should not be used to express a desire for the pod to be scheduled on a specific node. https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodename |
| `nodeSelector` | `object` | NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/ |
| `os` | `object` | PodOS defines the OS parameters of a pod. |
| `overhead` | `object` | Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md |
| `preemptionPolicy` | `string` | PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.<br><br>Possible enum values:<br>• `"Never"` means that pod never preempts other pods with lower priority.<br>• `"PreemptLowerPriority"` means that pod can preempt other pods with lower priority. |
| `priority` | `integer` | The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority. |

| Property | Type | Description |
| --- | --- | --- |
| `priorityClassName` | `string` | If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default. |
| `readinessGates` | `array` | If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates ↗ |
| `resourceClaims` | `array` | ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.<br><br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy ↗<br><br>Possible enum values:<br>• `"Always"`<br>• `"Never"`<br>• `"OnFailure"` |
| `runtimeClassName` | `string` | RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class ↗ |
| `schedulerName` | `string` | If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler. |
| `schedulingGates` | `array` | SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.<br><br>SchedulingGates can only be set at pod creation time, and be removed only afterwards. |

| Property | Type | Description |
| --- | --- | --- |
| `securityContext` | `object` | PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext. |
| `serviceAccount` | `string` | DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead. |
| `serviceAccountName` | `string` | ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/ ↗ |
| `setHostnameAsFQDN` | `boolean` | If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false. |
| `shareProcessNamespace` | `boolean` | Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false. |
| `subdomain` | `string` | If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all. |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds. |
| `tolerations` | `array` | If specified, the pod's tolerations. |
| `topologySpreadConstraints` | `array` | TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed. |

| Property | Type | Description |
| --- | --- | --- |
| `volumes` | `array` | List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes ↗ |

## .spec.jobTemplate.spec.template.spec.affinity

### Description

Affinity is a group of affinity scheduling rules.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `nodeAffinity` | `object` | Node affinity is a group of node affinity scheduling rules. |
| `podAffinity` | `object` | Pod affinity is a group of inter pod affinity scheduling rules. |
| `podAntiAffinity` | `object` | Pod anti affinity is a group of inter pod anti affinity scheduling rules. |

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity

### Description

Node affinity is a group of node affinity scheduling rules.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `object` | A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms. |

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

**Type**

`object`

**Required**

`weight` `preference`

| Property | Type | Description |
|---|---|---|
| `preference` | `object` | A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm. |
| `weight` | `integer` | Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100. |

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

**Type**

`object`

**Required**

`nodeSelectorTerms`

| Property | Type | Description |
| --- | --- | --- |
| `nodeSelectorTerms` | `array` | Required. A list of node selector terms. The terms are ORed. |

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms

**Description**

Required. A list of node selector terms. The terms are ORed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |

| Property | Type | Description |
| --- | --- | --- |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity

**Description**

Pod affinity is a group of inter pod affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the |

| Property | Type | Description |
| --- | --- | --- |
| | | corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
| --- | --- | --- |
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into |

| Property | Type | Description |
| --- | --- | --- |
| | | consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity

**Description**

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the |

| Property | Type | Description |
| --- | --- | --- |
| | | elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
| --- | --- | --- |
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

# .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.containers

**Description**

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[]

**Description**

A single application container that you want to run within a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |

| Property | Type | Description |
| --- | --- | --- |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |

| Property | Type | Description |
| --- | --- | --- |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.jobTemplate.spec.template.spec.containers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].args[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.containers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.containers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.jobTemplate.spec.template.spec.containers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.jobTemplate.spec.template.spec.containers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.jobTemplate.spec.template.spec.containers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.jobTemplate.spec.template.spec.containers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.jobTemplate.spec.template.spec.containers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.jobTemplate.spec.template.spec.containers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.jobTemplate.spec.template.spec.containers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.jobTemplate.spec.template.spec.containers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.postStart.exec.command[]

**Type**

`string`

# .spec.jobTemplate.spec.template.spec.containers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

# .spec.jobTemplate.spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

# .spec.jobTemplate.spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
|---|---|---|
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.jobTemplate.spec.template.spec.containers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.containers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.jobTemplate.spec.template.spec.containers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.containers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.containers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.jobTemplate.spec.template.spec.containers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.containers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.containers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.jobTemplate.spec.template.spec.containers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.jobTemplate.spec.template.spec.containers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.containers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.containers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.jobTemplate.spec.template.spec.containers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.containers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.containers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.jobTemplate.spec.template.spec.containers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.jobTemplate.spec.template.spec.containers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.jobTemplate.spec.template.spec.containers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.containers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.containers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.jobTemplate.spec.template.spec.containers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.jobTemplate.spec.template.spec.containers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.jobTemplate.spec.template.spec.containers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.containers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.containers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.jobTemplate.spec.template.spec.containers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.jobTemplate.spec.template.spec.containers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.jobTemplate.spec.template.spec.containers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.jobTemplate.spec.template.spec.containers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.containers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.jobTemplate.spec.template.spec.containers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.containers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.jobTemplate.spec.template.spec.containers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.containers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.containers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.containers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.jobTemplate.spec.template.spec.containers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.containers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.jobTemplate.spec.template.spec.dnsConfig

**Description**

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `nameservers` | `array` | A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed. |
| `options` | `array` | A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy. |
| `searches` | `array` | A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed. |

## .spec.jobTemplate.spec.template.spec.dnsConfig.nameservers

**Description**

A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.dnsConfig.nameservers[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.dnsConfig.options[]

**Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is this DNS resolver option's name. Required. |
| `value` | `string` | Value is this DNS resolver option's value. |

## .spec.jobTemplate.spec.template.spec.dnsConfig.searches

**Description**

A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.dnsConfig.searches[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers

**Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[]

**Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable |

| Property | Type | Description |
| --- | --- | --- |
| | | exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | Ports are not allowed for ephemeral containers. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for the container to manage the restart behavior of each container within a pod. This may only be set for init containers. You cannot set this field on ephemeral containers. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `targetContainerName` | `string` | If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.<br><br>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined. |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message |

| Property | Type | Description |
| --- | --- | --- |
| | | length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>- `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>- `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].args

**Description**

Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].args[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef

### Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

### Type

`object`

### Required

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

### Description

SecretKeySelector selects a key of a Secret.

### Type

`object`

### Required

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty.<br>Instances of this type with an empty value here are almost certainly wrong. More info: |

| Property | Type | Description |
| --- | --- | --- |
| | | https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
|---|---|---|
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].livenessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling |

| Property | Type | Description |
| --- | --- | --- |
| | | ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].ports

**Description**

Ports are not allowed for ephemeral containers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
|---|---|---|
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>- `"SCTP"` is the SCTP protocol.<br>- `"TCP"` is the TCP protocol.<br>- `"UDP"` is the UDP protocol. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].readinessProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br><br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |

| Property | Type | Description |
| --- | --- | --- |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br>• `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>• `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].startupProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling |

| Property | Type | Description |
| --- | --- | --- |
| | | ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].startupProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].startupProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
|---|---|---|
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.ephemeralContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
|---|---|---|
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other |

| Property | Type | Description |
| --- | --- | --- |
| | | containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br>If ReadOnly is false, this field has no meaning and must be unspecified.<br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.jobTemplate.spec.template.spec.hostAliases

**Description**

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.hostAliases[]

**Description**

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

**Type**

`object`

**Required**

`ip`

| Property | Type | Description |
| --- | --- | --- |
| `hostnames` | `array` | Hostnames for the above IP address. |
| `ip` | `string` | IP address of the host file entry. |

## .spec.jobTemplate.spec.template.spec.hostAliases[].hostnames

**Description**

Hostnames for the above IP address.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.hostAliases[].hostnames[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.imagePullSecrets

**Description**

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.imagePullSecrets[]

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.jobTemplate.spec.template.spec.initContainers

**Description**

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[]

**Description**

A single application container that you want to run within a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload |

| Property | Type | Description |
| --- | --- | --- |
| | | controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers |

| Property | Type | Description |
| --- | --- | --- |
| | | have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |

| Property | Type | Description |
| --- | --- | --- |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.jobTemplate.spec.template.spec.initContainers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].args[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.initContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.initContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.jobTemplate.spec.template.spec.initContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.jobTemplate.spec.template.spec.initContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.jobTemplate.spec.template.spec.initContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.jobTemplate.spec.template.spec.initContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.jobTemplate.spec.template.spec.initContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.jobTemplate.spec.template.spec.initContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.jobTemplate.spec.template.spec.initContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.jobTemplate.spec.template.spec.initContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
|---|---|---|
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br><br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
|---|---|---|
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.jobTemplate.spec.template.spec.initContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.initContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.jobTemplate.spec.template.spec.initContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.initContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.initContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.jobTemplate.spec.template.spec.initContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.initContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.initContainers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.jobTemplate.spec.template.spec.initContainers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.jobTemplate.spec.template.spec.initContainers[].readinessProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.initContainers[].readinessProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.initContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.jobTemplate.spec.template.spec.initContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.initContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.initContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.jobTemplate.spec.template.spec.initContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.jobTemplate.spec.template.spec.initContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.jobTemplate.spec.template.spec.initContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.initContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.initContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>- `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>- `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.jobTemplate.spec.template.spec.initContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.jobTemplate.spec.template.spec.initContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.jobTemplate.spec.template.spec.initContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.initContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.initContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.jobTemplate.spec.template.spec.initContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.jobTemplate.spec.template.spec.initContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.jobTemplate.spec.template.spec.initContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.jobTemplate.spec.template.spec.initContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.jobTemplate.spec.template.spec.initContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.initContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.jobTemplate.spec.template.spec.initContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.jobTemplate.spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.jobTemplate.spec.template.spec.initContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.jobTemplate.spec.template.spec.initContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.jobTemplate.spec.template.spec.initContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.initContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.jobTemplate.spec.template.spec.nodeSelector

**Description**

NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.os

**Description**

PodOS defines the OS parameters of a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration ↗ Clients should expect to handle additional values and treat unrecognized values in this field as os: null |

## .spec.jobTemplate.spec.template.spec.overhead

**Description**

Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.readinessGates

**Description**

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.readinessGates[]

**Description**

PodReadinessGate contains the reference to a pod condition

**Type**

`object`

**Required**

`conditionType`

| Property | Type | Description |
| --- | --- | --- |
| `conditionType` | `string` | ConditionType refers to a condition in the pod's condition list with matching type. |

## .spec.jobTemplate.spec.template.spec.resourceClaims

**Description**

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.resourceClaims[]

**Description**

PodResourceClaim references exactly one ResourceClaim, either directly or by naming a ResourceClaimTemplate which is then turned into a ResourceClaim for the pod. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL. |
| `resourceClaimName` | `string` | ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |

| Property | Type | Description |
| --- | --- | --- |
| `resourceClaimTemplateName` | `string` | ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.<br><br>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The pod name and resource name, along with a generated component, will be used to form a unique name for the ResourceClaim, which will be recorded in pod.status.resourceClaimStatuses.<br><br>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |

## .spec.jobTemplate.spec.template.spec.resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br><br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.jobTemplate.spec.template.spec.resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.jobTemplate.spec.template.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.schedulingGates

**Description**

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod. SchedulingGates can only be set at pod creation time, and be removed only afterwards.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.schedulingGates[]

**Description**

PodSchedulingGate is associated to a Pod to guard its scheduling.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the scheduling gate. Each scheduling gate must have a unique name field. |

## .spec.jobTemplate.spec.template.spec.securityContext

**Description**

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `fsGroup` | `integer` | A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:<br>1. The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw----<br>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows. |
| `fsGroupChangePolicy` | `string` | fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership(and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.<br>Possible enum values:<br>- `"Always"` indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior.<br>- `"OnRootMismatch"` indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxChangePolicy` | `string` | seLinuxChangePolicy defines how the container's SELinux label is applied to all volumes used by the Pod. It has no effect on nodes that do not support SELinux or to volumes does not support SELinux. Valid values are "MountOption" and "Recursive".<br><br>"Recursive" means relabeling of all files on all Pod volumes by the container runtime. This may be slow for large volumes, but allows mixing privileged and unprivileged Pods sharing the same volume on the same node.<br><br>"MountOption" mounts all eligible Pod volumes with `-o context` mount option. This requires all Pods that share the same volume to use the same SELinux label. It is not possible to share the same volume among privileged and unprivileged Pods. Eligible volumes are in-tree FibreChannel and iSCSI volumes, and all CSI volumes whose CSI driver announces SELinux support by setting spec.seLinuxMount: true in their CSIDriver instance. Other volumes are always re-labelled recursively. "MountOption" value is allowed only when SELinuxMount feature gate is enabled.<br><br>If not specified and SELinuxMount feature gate is enabled, "MountOption" is used. If not specified and SELinuxMount feature gate is disabled, "MountOption" is used for ReadWriteOncePod volumes and "Recursive" for all other volumes.<br><br>This field affects only Pods that have SELinux label set, either in PodSecurityContext or in SecurityContext of all containers.<br><br>All Pods that use the same volume should use the same seLinuxChangePolicy, otherwise some pods can get stuck in ContainerCreating state. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |

| Property | Type | Description |
| --- | --- | --- |
| `supplementalGroups` | `array` | A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows. |
| `supplementalGroupsPolicy` | `string` | Defines how supplemental groups of the first container processes are calculated. Valid values are "Merge" and "Strict". If not specified, "Merge" is used. (Alpha) Using the field requires the SupplementalGroupsPolicy feature gate to be enabled and the container runtime must implement support for this feature. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>- `"Merge"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be merged with the primary user's groups as defined in the container image (in /etc/group).<br>- `"Strict"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be used instead of any groups defined in the container image. |
| `sysctls` | `array` | Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.jobTemplate.spec.template.spec.securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values: |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>• `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.jobTemplate.spec.template.spec.securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.jobTemplate.spec.template.spec.securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values: |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.jobTemplate.spec.template.spec.securityContext.supplementalGroups

**Description**

A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.securityContext.supplementalGroups[]

**Type**

`integer`

## .spec.jobTemplate.spec.template.spec.securityContext.sysctls

**Description**

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.securityContext.sysctls[]

**Description**

Sysctl defines a kernel parameter to be set

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of a property to set |
| `value` | `string` | Value of a property to set |

## .spec.jobTemplate.spec.template.spec.securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.jobTemplate.spec.template.spec.tolerations

**Description**

If specified, the pod's tolerations.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.tolerations[]

**Description**

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `effect` | `string` | Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.<br><br>Possible enum values: |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"NoExecute"` Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController.<br>• `"NoSchedule"` Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler.<br>• `"PreferNoSchedule"` Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler. |
| `key` | `string` | Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys. |
| `operator` | `string` | Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.<br>Possible enum values:<br>• `"Equal"`<br>• `"Exists"` |
| `tolerationSeconds` | `integer` | TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system. |
| `value` | `string` | Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string. |

## .spec.jobTemplate.spec.template.spec.topologySpreadConstraints

**Description**

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.topologySpreadConstraints[]

**Description**

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

**Type**

`object`

**Required**

`maxSkew` `topologyKey` `whenUnsatisfiable`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.<br><br>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default). |
| `maxSkew` | `integer` | MaxSkew describes the degree to which pods may be unevenly distributed. When `whenUnsatisfiable=DoNotSchedule`, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1. \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P \| - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When `whenUnsatisfiable=ScheduleAnyway`, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed. |
| `minDomains` | `integer` | MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.<br><br>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2: \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P P \| The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew. |
| `nodeAffinityPolicy` | `string` | NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.<br><br>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br><br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |
| `nodeTaintsPolicy` | `string` | NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.<br><br>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br><br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |
| `topologyKey` | `string` | TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field. |
| `whenUnsatisfiable` | `string` | WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: \| zone1 \| zone2 \| zone3 \| \| P P P \| P \| P \| If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it *more* imbalanced. It's a required field.<br><br>Possible enum values:<br><br>• `"DoNotSchedule"` instructs the scheduler not to schedule the pod when constraints are not satisfied.<br>• `"ScheduleAnyway"` instructs the scheduler to schedule the pod even if constraints are not satisfied. |

## .spec.jobTemplate.spec.template.spec.topologySpreadConstraints[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.topologySpreadConstraints[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.topologySpreadConstraints[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector. This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.topologySpreadConstraints[].matchLabelKeys[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.volumes

**Description**

List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[]

**Description**

Volume represents a named volume in a pod that may be accessed by any container in the pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `awsElasticBlockStore` | `object` | Represents a Persistent Disk resource in AWS.<br><br>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling. |
| `azureDisk` | `object` | AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod. |
| `azureFile` | `object` | AzureFile represents an Azure File Service mount on the host and bind mount to the pod. |
| `cephfs` | `object` | Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling. |
| `cinder` | `object` | Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling. |
| `configMap` | `object` | Adapts a ConfigMap into a volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling. |
| `csi` | `object` | Represents a source location of a volume to mount, managed by an external CSI driver |
| `downwardAPI` | `object` | DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling. |
| `emptyDir` | `object` | Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling. |
| `ephemeral` | `object` | Represents an ephemeral volume that is handled by a normal storage driver. |
| `fc` | `object` | Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling. |
| `flexVolume` | `object` | FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin. |

| Property | Type | Description |
| --- | --- | --- |
| `flocker` | `object` | Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling. |
| `gcePersistentDisk` | `object` | Represents a Persistent Disk resource in Google Compute Engine.<br><br>A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling. |
| `gitRepo` | `object` | Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.<br><br>DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container. |
| `glusterfs` | `object` | Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling. |
| `hostPath` | `object` | Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling. |
| `image` | `object` | ImageVolumeSource represents a image volume resource. |
| `iscsi` | `object` | Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling. |
| `name` | `string` | name of the volume. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `nfs` | `object` | Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling. |
| `persistentVolumeClaim` | `object` | PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system). |
| `photonPersistentDisk` | `object` | Represents a Photon Controller persistent disk resource. |

| Property | Type | Description |
| --- | --- | --- |
| `portworxVolume` | `object` | PortworxVolumeSource represents a Portworx volume resource. |
| `projected` | `object` | Represents a projected volume source |
| `quobyte` | `object` | Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling. |
| `rbd` | `object` | Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling. |
| `scaleIO` | `object` | ScaleIOVolumeSource represents a persistent ScaleIO volume |
| `secret` | `object` | Adapts a Secret into a volume.<br><br>The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling. |
| `storageos` | `object` | Represents a StorageOS persistent volume resource. |
| `vsphereVolume` | `object` | Represents a vSphere volume resource. |

## .spec.jobTemplate.spec.template.spec.volumes[].awsElasticBlockStore

### Description

Represents a Persistent Disk resource in AWS. An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). |
| `readOnly` | `boolean` | readOnly value true will force the readOnly setting in VolumeMounts. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `volumeID` | `string` | volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].azureDisk

### Description

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

### Type

`object`

### Required

`diskName` `diskURI`

| Property | Type | Description |
| --- | --- | --- |
| `cachingMode` | `string` | cachingMode is the Host Caching mode: None, Read Only, Read Write.<br><br>Possible enum values:<br>• `"None"`<br>• `"ReadOnly"`<br>• `"ReadWrite"` |
| `diskName` | `string` | diskName is the Name of the data disk in the blob storage |
| `diskURI` | `string` | diskURI is the URI of data disk in the blob storage |
| `fsType` | `string` | fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `kind` | `string` | kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared<br><br>Possible enum values: |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Dedicated"`<br>• `"Managed"`<br>• `"Shared"` |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

## .spec.jobTemplate.spec.template.spec.volumes[].azureFile

### Description

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

### Type

`object`

### Required

`secretName` `shareName`

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretName` | `string` | secretName is the name of secret that contains Azure Storage Account Name and Key |
| `shareName` | `string` | shareName is the azure share Name |

## .spec.jobTemplate.spec.template.spec.volumes[].cephfs

### Description

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`monitors`

| Property | Type | Description |
| --- | --- | --- |
| `monitors` | `array` | monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `path` | `string` | path is Optional: Used as the mounted root, rather than the full Ceph tree, default is / |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it |
| `secretFile` | `string` | secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it |

## .spec.jobTemplate.spec.template.spec.volumes[].cephfs.monitors

**Description**

monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].cephfs.monitors[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.volumes[].cephfs.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |

## .spec.jobTemplate.spec.template.spec.volumes[].cinder

**Description**

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeID` | `string` | volumeID used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].cinder.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].configMap

**Description**

Adapts a ConfigMap into a volume. The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires |

| Property | Type | Description |
| --- | --- | --- |
| | | decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.jobTemplate.spec.template.spec.volumes[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.jobTemplate.spec.template.spec.volumes[].csi

**Description**

Represents a source location of a volume to mount, managed by an external CSI driver

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
| --- | --- | --- |
| `driver` | `string` | driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster. |
| `fsType` | `string` | fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply. |
| `nodePublishSecretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `readOnly` | `boolean` | readOnly specifies a read-only configuration for the volume. Defaults to false (read/write). |
| `volumeAttributes` | `object` | volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values. |

## .spec.jobTemplate.spec.template.spec.volumes[].csi.nodePublishSecretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].csi.volumeAttributes

**Description**

volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.volumes[].downwardAPI

**Description**

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | Items is a list of downward API volume file |

## .spec.jobTemplate.spec.template.spec.volumes[].downwardAPI.items

**Description**

Items is a list of downward API volume file

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |

| Property | Type | Description |
| --- | --- | --- |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.jobTemplate.spec.template.spec.volumes[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.jobTemplate.spec.template.spec.volumes[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and |

| Property | Type | Description |
| --- | --- | --- |
| | | The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.jobTemplate.spec.template.spec.volumes[].emptyDir

### Description

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `medium` | `string` | medium represents what type of storage medium should back this directory. The default is "" which means to use the n |
| `sizeLimit` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAM<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits> <r<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.html)<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a number<br>When a Quantity is parsed from a string, it will remember the type of suffix it had, an<br>Before serializing, Quantity will be put in "canonical form". This means that Exponent/<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suffix<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point number.<br>Non-canonical values will still parse as long as they are well formed, but will be re-e<br>This format is intended to make it difficult to use these numbers without writing some |

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral

**Description**

Represents an ephemeral volume that is handled by a normal storage driver.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `volumeClaimTemplate` | `object` | PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource. |

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate

**Description**

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

**Type**

`object`

**Required**

`spec`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes |

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec

### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `accessModes` | `array` | accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1 ↗ |
| `dataSource` | `object` | TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace. |
| `dataSourceRef` | `object` | TypedObjectReference contains enough information to let you locate the typed referenced object |
| `resources` | `object` | VolumeResourceRequirements describes the storage resource requirements for a volume. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `storageClassName` | `string` | storageClassName is the name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1 ↗ |
| `volumeAttributesClassName` | `string` | volumeAttributesClassName may be used to set the VolumeAttributesClass used by this claim. If specified, the CSI driver will create or update the volume with the attributes defined in the corresponding VolumeAttributesClass. This has a different purpose than storageClassName, it can be changed after the claim is created. An empty string value means that no VolumeAttributesClass will be applied to the claim but it's not allowed to reset this field to empty string once it is set. If unspecified and the PersistentVolumeClaim is unbound, the default VolumeAttributesClass will be set by the persistentvolume controller if it exists. If the resource referred to by volumeAttributesClass does not exist, this PersistentVolumeClaim will be set to a Pending state, as reflected by the |

| Property | Type | Description |
| --- | --- | --- |
| | | modifyVolumeStatus field, until such as a resource exists. More info: https://kubernetes.io/docs/concepts/storage/volume-attributes-classes/ (Beta) Using this field requires the VolumeAttributesClass feature gate to be enabled (off by default). |
| `volumeMode` | `string` | volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.<br><br>Possible enum values:<br>• `"Block"` means the volume will not be formatted with a filesystem and will remain a raw block device.<br>• `"Filesystem"` means the volume will be or is formatted with a filesystem. |
| `volumeName` | `string` | volumeName is the binding reference to the PersistentVolume backing this claim. |

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes

**Description**

accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

**Description**

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef

**Description**

TypedObjectReference contains enough information to let you locate the typed referenced object

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |
| `namespace` | `string` | Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled. |

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

**Description**

VolumeResourceRequirements describes the storage resource requirements for a volume.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.volumes[].fc

**Description**

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `lun` | `integer` | lun is Optional: FC target lun number |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `targetWWNs` | `array` | targetWWNs is Optional: FC target worldwide names (WWNs) |
| `wwids` | `array` | wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously. |

## .spec.jobTemplate.spec.template.spec.volumes[].fc.targetWWNs

**Description**

targetWWNs is Optional: FC target worldwide names (WWNs)

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].fc.targetWWNs[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.volumes[].fc.wwids

**Description**

wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].fc.wwids[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.volumes[].flexVolume

**Description**

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
| --- | --- | --- |
| `driver` | `string` | driver is the name of the driver to use for this volume. |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script. |
| `options` | `object` | options is Optional: this field holds extra command options if any. |
| `readOnly` | `boolean` | readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |

## .spec.jobTemplate.spec.template.spec.volumes[].flexVolume.options

**Description**

options is Optional: this field holds extra command options if any.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.volumes[].flexVolume.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].flocker

**Description**

Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `datasetName` | `string` | datasetName is Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated |
| `datasetUUID` | `string` | datasetUUID is the UUID of the dataset. This is unique identifier of a Flocker dataset |

## .spec.jobTemplate.spec.template.spec.volumes[].gcePersistentDisk

**Description**

Represents a Persistent Disk resource in Google Compute Engine. A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`pdName`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or |

| Property | Type | Description |
| --- | --- | --- |
| | | you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `pdName` | `string` | pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].gitRepo

**Description**

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

**Type**

`object`

**Required**

`repository`

| Property | Type | Description |
| --- | --- | --- |
| `directory` | `string` | directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name. |
| `repository` | `string` | repository is the URL |
| `revision` | `string` | revision is the commit hash for the specified revision. |

## .spec.jobTemplate.spec.template.spec.volumes[].glusterfs

**Description**

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`endpoints` `path`

| Property | Type | Description |
| --- | --- | --- |
| `endpoints` | `string` | endpoints is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |
| `path` | `string` | path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |
| `readOnly` | `boolean` | readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].hostPath

**Description**

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗ |
| `type` | `string` | type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗<br><br>Possible enum values:<br>• `""` For backwards compatible, leave it empty if unset<br>• `"BlockDevice"` A block device must exist at the given path<br>• `"CharDevice"` A character device must exist at the given path<br>• `"Directory"` A directory must exist at the given path<br>• `"DirectoryOrCreate"` If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet.<br>• `"File"` A file must exist at the given path<br>• `"FileOrCreate"` If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet.<br>• `"Socket"` A UNIX socket must exist at the given path |

## .spec.jobTemplate.spec.template.spec.volumes[].image

**Description**

ImageVolumeSource represents a image volume resource.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `pullPolicy` | `string` | Policy for pulling OCI objects. Possible values are: Always: the kubelet always attempts to pull the reference. Container creation will fail If the pull fails. Never: the kubelet never pulls the reference and only uses a local image or artifact. Container creation will fail if the reference isn't present. IfNotPresent: the kubelet pulls if the reference isn't already present on disk. Container creation will fail if the reference isn't present and the pull fails. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise.<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `reference` | `string` | Required: Image or artifact reference to be used. Behaves in the same way as pod.spec.containers[*].image. Pull secrets will be assembled in the same way as for the container image by looking up node credentials, SA image pull secrets, and pod spec image pull secrets. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |

## .spec.jobTemplate.spec.template.spec.volumes[].iscsi

**Description**

Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`targetPortal` `iqn` `lun`

| Property | Type | Description |
| --- | --- | --- |
| `chapAuthDiscovery` | `boolean` | chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication |
| `chapAuthSession` | `boolean` | chapAuthSession defines whether support iSCSI Session CHAP authentication |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi |

| Property | Type | Description |
| --- | --- | --- |
| `initiatorName` | `string` | initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection. |
| `iqn` | `string` | iqn is the target iSCSI Qualified Name. |
| `iscsiInterface` | `string` | iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp). |
| `lun` | `integer` | lun represents iSCSI Target Lun number. |
| `portals` | `array` | portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `targetPortal` | `string` | targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |

## .spec.jobTemplate.spec.template.spec.volumes[].iscsi.portals

**Description**

portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].iscsi.portals[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.volumes[].iscsi.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].nfs

### Description

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`server` `path`

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs ↗ |
| `readOnly` | `boolean` | readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs ↗ |
| `server` | `string` | server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].persistentVolumeClaim

### Description

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

### Type

`object`

### Required

`claimName`

| Property | Type | Description |
| --- | --- | --- |
| `claimName` | `string` | claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims ↗ |
| `readOnly` | `boolean` | readOnly Will force the ReadOnly setting in VolumeMounts. Default false. |

## .spec.jobTemplate.spec.template.spec.volumes[].photonPersistentDisk

**Description**

Represents a Photon Controller persistent disk resource.

**Type**

`object`

**Required**

`pdID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `pdID` | `string` | pdID is the ID that identifies Photon Controller persistent disk |

## .spec.jobTemplate.spec.template.spec.volumes[].portworxVolume

**Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `volumeID` | `string` | volumeID uniquely identifies a Portworx volume |

## .spec.jobTemplate.spec.template.spec.volumes[].projected

**Description**

Represents a projected volume source

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `sources` | `array` | sources is the list of volume projections. Each entry in this list handles one source. |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources

**Description**

sources is the list of volume projections. Each entry in this list handles one source.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[]

**Description**

Projection that may be projected along with other supported volume types. Exactly one of these fields must be set.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `clusterTrustBundle` | `object` | ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem. |
| `configMap` | `object` | Adapts a ConfigMap into a projected volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode. |
| `downwardAPI` | `object` | Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode. |
| `secret` | `object` | Adapts a secret into a projected volume.<br><br>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode. |
| `serviceAccountToken` | `object` | ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API |

| Property | Type | Description |
| --- | --- | --- |
| | | Server or otherwise). |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].clusterTrustBundle

**Description**

ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem.

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `name` | `string` | Select a single ClusterTrustBundle by object name. Mutually-exclusive with signerName and labelSelector. |
| `optional` | `boolean` | If true, don't block pod startup if the referenced ClusterTrustBundle(s) aren't available. If using name, then the named ClusterTrustBundle is allowed not to exist. If using signerName, then the combination of signerName and labelSelector is allowed to match zero ClusterTrustBundles. |
| `path` | `string` | Relative path from the volume root to write the bundle. |
| `signerName` | `string` | Select all ClusterTrustBundles that match this signer name. Mutually-exclusive with name. The contents of all selected ClusterTrustBundles will be unified and deduplicated. |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |

| Property | Type | Description |
| --- | --- | --- |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].configMap

**Description**

Adapts a ConfigMap into a projected volume. The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].downwardAPI

**Description**

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `items` | `array` | Items is a list of DownwardAPIVolume file |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].downwardAPI.items

**Description**

Items is a list of DownwardAPIVolume file

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
|---|---|---|
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].secret

**Description**

Adapts a secret into a projected volume. The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional field specify whether the Secret or its key must be defined |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].secret.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.jobTemplate.spec.template.spec.volumes[].projected.sources[].serviceAccountToken

### Description

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

### Type

`object`

### Required

`path`

| Property | Type | Description |
| --- | --- | --- |
| `audience` | `string` | audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver. |
| `expirationSeconds` | `integer` | expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours.Defaults to 1 hour and must be at least 10 minutes. |
| `path` | `string` | path is the path relative to the mount point of the file to project the token into. |

## .spec.jobTemplate.spec.template.spec.volumes[].quobyte

### Description

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`registry` `volume`

| Property | Type | Description |
| --- | --- | --- |
| `group` | `string` | group to map volume access to Default is no group |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false. |
| `registry` | `string` | registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes |
| `tenant` | `string` | tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin |
| `user` | `string` | user to map volume access to Defaults to serivceaccount user |
| `volume` | `string` | volume is a string that references an already created Quobyte volume by name. |

## .spec.jobTemplate.spec.template.spec.volumes[].rbd

**Description**

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`monitors` `image`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd ↗ |
| `image` | `string` | image is the rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `keyring` | `string` | keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `monitors` | `array` | monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `pool` | `string` | pool is the rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is the rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].rbd.monitors

**Description**

monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].rbd.monitors[]

**Type**

`string`

## .spec.jobTemplate.spec.template.spec.volumes[].rbd.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].scaleIO

**Description**

ScaleIOVolumeSource represents a persistent ScaleIO volume

**Type**

`object`

**Required**

`gateway` `system` `secretRef`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs". |
| `gateway` | `string` | gateway is the host address of the ScaleIO API Gateway. |
| `protectionDomain` | `string` | protectionDomain is the name of the ScaleIO Protection Domain for the configured storage. |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `sslEnabled` | `boolean` | sslEnabled Flag enable/disable SSL communication with Gateway, default false |
| `storageMode` | `string` | storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned. |
| `storagePool` | `string` | storagePool is the ScaleIO Storage Pool associated with the protection domain. |
| `system` | `string` | system is the name of the storage system as configured in ScaleIO. |
| `volumeName` | `string` | volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source. |

## .spec.jobTemplate.spec.template.spec.volumes[].scaleIO.secretRef

### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].secret

**Description**

Adapts a Secret into a volume. The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `defaultMode` | `integer` | defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `optional` | `boolean` | optional field specify whether the Secret or its keys must be defined |
| `secretName` | `string` | secretName is the name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].secret.items

**Description**

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.jobTemplate.spec.template.spec.volumes[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.jobTemplate.spec.template.spec.volumes[].storageos

**Description**

Represents a StorageOS persistent volume resource.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeName` | `string` | volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace. |
| `volumeNamespace` | `string` | volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created. |

## .spec.jobTemplate.spec.template.spec.volumes[].storageos.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.jobTemplate.spec.template.spec.volumes[].vsphereVolume

### Description

Represents a vSphere volume resource.

### Type

`object`

### Required

`volumePath`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `storagePolicyID` | `string` | storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName. |
| `storagePolicyName` | `string` | storagePolicyName is the storage Policy Based Management (SPBM) profile name. |
| `volumePath` | `string` | volumePath is the path that identifies vSphere volume vmdk |

## .status

### Description

CronJobStatus represents the current state of a cron job.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `active` | `array` | A list of pointers to currently running jobs. |
| `lastScheduleTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

| Property | Type | Description |
| --- | --- | --- |
| `lastSuccessfulTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.active

### Description

A list of pointers to currently running jobs.

### Type

`array`

## .status.active[]

### Description

ObjectReference contains enough information to let you inspect or modify the referred object.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | API version of the referent. |
| `fieldPath` | `string` | If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as desiredState.manifest.containers[2]. For example, if the object reference is to a container within a pod, this would take on a value like: "spec.containers{name}" (where "name" refers to the name of the container that triggered the event) or if no container name is specified "spec.containers[2]" (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object. |
| `kind` | `string` | Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds ↗ |
| `name` | `string` | Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `namespace` | `string` | Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/ ↗ |
| `resourceVersion` | `string` | Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `uid` | `string` | UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids ↗ |

# API Endpoints

The following API endpoints are available:

- `/kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/cronjobs`
  - `DELETE` : delete collection of CronJob
  - `GET` : list objects of kind CronJob
  - `POST` : create a new CronJob
- `/kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/cronjobs/{name}`
  - `DELETE` : delete the specified CronJob
  - `GET` : read the specified CronJob
  - `PATCH` : partially update the specified CronJob
  - `PUT` : replace the specified CronJob
- `/kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/cronjobs/{name}/status`
  - `GET` : read status of the specified CronJob
  - `PATCH` : partially update status of the specified CronJob
  - `PUT` : replace status of the specified CronJob

## /kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/cronjobs

### HTTP method

`DELETE`

### Description

delete collection of CronJob

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

### Description

list objects of kind CronJob

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `CronJobList` schema |

| HTTP code | Response body |
| --- | --- |
| 401 - Unauthorized | Empty |

### HTTP method

`POST`

### Description

create a new CronJob

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `CronJob` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `CronJob` schema |
| 201 - Created | `CronJob` schema |
| 202 - Accepted | `CronJob` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/cronjobs/{name}

### HTTP method

`DELETE`

### Description

delete the specified CronJob

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 202 - Accepted | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

### Description

read the specified CronJob

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `CronJob` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update the specified CronJob

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `CronJob` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace the specified CronJob

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
|---|---|---|
| `body` | `CronJob` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `CronJob` schema |
| 201 - Created | `CronJob` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/cronjobs/{name}/status

### HTTP method

`GET`

### Description

read status of the specified CronJob

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `CronJob` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update status of the specified CronJob

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `CronJob` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace status of the specified CronJob

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
|---|---|---|
| `body` | `CronJob` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `CronJob` schema |
| 201 - Created | `CronJob` schema |
| 401 - Unauthorized | Empty |

# DaemonSet [apps/v1]

**Description**

DaemonSet represents the configuration of a daemon set.

**Type**

`object`

## Specification

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources |
| `kind` | `string` | Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | DaemonSetSpec is the specification of a daemon set. |
| `status` | `object` | DaemonSetStatus represents the current status of a daemon set. |

### .spec

**Description**

DaemonSetSpec is the specification of a daemon set.

**Type**

`object`

**Required**

`selector` `template`

| Property | Type | Description |
| --- | --- | --- |
| `minReadySeconds` | `integer` | The minimum number of seconds for which a newly created DaemonSet pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready). |
| `revisionHistoryLimit` | `integer` | The number of old history to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `template` | `object` | PodTemplateSpec describes the data a pod should have when created from a template |
| `updateStrategy` | `object` | DaemonSetUpdateStrategy is a struct used to control the update strategy for a DaemonSet. |

## .spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template

**Description**

PodTemplateSpec describes the data a pod should have when created from a template

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PodSpec is a description of a pod. |

## .spec.template.spec

### Description

PodSpec is a description of a pod.

### Type

`object`

### Required

`containers`

| Property | Type | Description |
| --- | --- | --- |
| `activeDeadlineSeconds` | `integer` | Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer. |
| `affinity` | `object` | Affinity is a group of affinity scheduling rules. |
| `automountServiceAccountToken` | `boolean` | AutomountServiceAccountToken indicates whether a service account token should be automatically mounted. |
| `containers` | `array` | List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated. |
| `dnsConfig` | `object` | PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy. |
| `dnsPolicy` | `string` | Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.<br><br>Possible enum values:<br>- `"ClusterFirst"` indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>- `"ClusterFirstWithHostNet"` indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Default"` indicates that the pod should use the default (as determined by kubelet) DNS settings.<br>• `"None"` indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig. |
| `enableServiceLinks` | `boolean` | EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching traditional container linking syntax. Optional: Defaults to true. |
| `ephemeralContainers` | `array` | List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. |
| `hostAliases` | `array` | HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. |
| `hostIPC` | `boolean` | Use the host's ipc namespace. Optional: Default to false. |
| `hostNetwork` | `boolean` | Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false. |
| `hostPID` | `boolean` | Use the host's pid namespace. Optional: Default to false. |
| `hostUsers` | `boolean` | Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new userns is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature. |
| `hostname` | `string` | Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value. |
| `imagePullSecrets` | `array` | ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `initContainers` | `array` | List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/ |
| `nodeName` | `string` | NodeName indicates in which node this pod is scheduled. If empty, this pod is a candidate for scheduling by the scheduler defined in schedulerName. Once this field is set, the kubelet for this node becomes responsible for the lifecycle of this pod. This field should not be used to express a desire for the pod to be scheduled on a specific node. https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodename |
| `nodeSelector` | `object` | NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/ |
| `os` | `object` | PodOS defines the OS parameters of a pod. |
| `overhead` | `object` | Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md |
| `preemptionPolicy` | `string` | PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.<br><br>Possible enum values:<br>• `"Never"` means that pod never preempts other pods with lower priority.<br>• `"PreemptLowerPriority"` means that pod can preempt other pods with lower priority. |
| `priority` | `integer` | The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority. |

| Property | Type | Description |
| --- | --- | --- |
| `priorityClassName` | `string` | If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default. |
| `readinessGates` | `array` | If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates ↗ |
| `resourceClaims` | `array` | ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.<br><br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy ↗<br><br>Possible enum values:<br>• `"Always"`<br>• `"Never"`<br>• `"OnFailure"` |
| `runtimeClassName` | `string` | RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class ↗ |
| `schedulerName` | `string` | If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler. |
| `schedulingGates` | `array` | SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.<br><br>SchedulingGates can only be set at pod creation time, and be removed only afterwards. |

| Property | Type | Description |
| --- | --- | --- |
| `securityContext` | `object` | PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext. |
| `serviceAccount` | `string` | DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead. |
| `serviceAccountName` | `string` | ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/ ↗ |
| `setHostnameAsFQDN` | `boolean` | If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false. |
| `shareProcessNamespace` | `boolean` | Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false. |
| `subdomain` | `string` | If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all. |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds. |
| `tolerations` | `array` | If specified, the pod's tolerations. |
| `topologySpreadConstraints` | `array` | TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed. |

| Property | Type | Description |
| --- | --- | --- |
| `volumes` | `array` | List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes ↗ |

## .spec.template.spec.affinity

**Description**

Affinity is a group of affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `nodeAffinity` | `object` | Node affinity is a group of node affinity scheduling rules. |
| `podAffinity` | `object` | Pod affinity is a group of inter pod affinity scheduling rules. |
| `podAntiAffinity` | `object` | Pod anti affinity is a group of inter pod anti affinity scheduling rules. |

## .spec.template.spec.affinity.nodeAffinity

**Description**

Node affinity is a group of node affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `object` | A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms. |

# .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

**Type**

`array`

# .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

**Type**

`object`

**Required**

`weight` `preference`

| Property | Type | Description |
|---|---|---|
| `preference` | `object` | A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm. |
| `weight` | `integer` | Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100. |

# .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

**Type**

`object`

**Required**

`nodeSelectorTerms`

| Property | Type | Description |
| --- | --- | --- |
| `nodeSelectorTerms` | `array` | Required. A list of node selector terms. The terms are ORed. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms

**Description**

Required. A list of node selector terms. The terms are ORed.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |

| Property | Type | Description |
| --- | --- | --- |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity

**Description**

Pod affinity is a group of inter pod affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the |

| Property | Type | Description |
| --- | --- | --- |
| | | corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
| --- | --- | --- |
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into |

| Property | Type | Description |
| --- | --- | --- |
| | | consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity

**Description**

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the |

| Property | Type | Description |
| --- | --- | --- |
| | | elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
| --- | --- | --- |
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

# .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into |

| Property | Type | Description |
| --- | --- | --- |
| | | consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.containers

**Description**

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[]

**Description**

A single application container that you want to run within a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |

| Property | Type | Description |
| --- | --- | --- |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |

| Property | Type | Description |
| --- | --- | --- |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.containers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].args[]

**Type**

`string`

## .spec.template.spec.containers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].command[]

**Type**

`string`

## .spec.template.spec.containers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.containers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.containers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.containers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.containers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.containers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.containers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.containers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.containers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.containers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.containers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.postStart.exec.command[]

**Type**

`string`

# .spec.template.spec.containers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

# .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

# .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.containers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.containers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.containers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.containers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.containers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.containers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.containers[].securityContext.appArmorProfile

### Description

AppArmorProfile defines a pod or container's AppArmor settings.

### Type

`object`

### Required

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.containers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.containers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.containers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.containers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.containers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.containers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.containers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.containers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.dnsConfig

**Description**

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `nameservers` | `array` | A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed. |
| `options` | `array` | A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy. |
| `searches` | `array` | A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed. |

## .spec.template.spec.dnsConfig.nameservers

**Description**

A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.nameservers[]

**Type**

`string`

## .spec.template.spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

`array`

## .spec.template.spec.dnsConfig.options[]

**Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is this DNS resolver option's name. Required. |
| `value` | `string` | Value is this DNS resolver option's value. |

## .spec.template.spec.dnsConfig.searches

**Description**

A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.searches[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers

**Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[]

**Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable |

| Property | Type | Description |
| --- | --- | --- |
| | | exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | Ports are not allowed for ephemeral containers. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for the container to manage the restart behavior of each container within a pod. This may only be set for init containers. You cannot set this field on ephemeral containers. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `targetContainerName` | `string` | If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.<br><br>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined. |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message |

| Property | Type | Description |
| --- | --- | --- |
| | | length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>- `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>- `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.ephemeralContainers[].args

**Description**

Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].args[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.ephemeralContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.ephemeralContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

# .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

# .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

# .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].ports

**Description**

Ports are not allowed for ephemeral containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |

| Property | Type | Description |
| --- | --- | --- |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br><br>Possible enum values:<br>- `"SCTP"` is the SCTP protocol.<br>- `"TCP"` is the TCP protocol.<br>- `"UDP"` is the UDP protocol. |

## .spec.template.spec.ephemeralContainers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.ephemeralContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.ephemeralContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.ephemeralContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.ephemeralContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.ephemeralContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.ephemeralContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.ephemeralContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.ephemeralContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.hostAliases

**Description**

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified.

**Type**

`array`

## .spec.template.spec.hostAliases[]

**Description**

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

**Type**

`object`

**Required**

`ip`

| Property | Type | Description |
|---|---|---|
| `hostnames` | `array` | Hostnames for the above IP address. |
| `ip` | `string` | IP address of the host file entry. |

## .spec.template.spec.hostAliases[].hostnames

**Description**

Hostnames for the above IP address.

**Type**

`array`

## .spec.template.spec.hostAliases[].hostnames[]

**Type**

`string`

## .spec.template.spec.imagePullSecrets

**Description**

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

**Type**

`array`

## .spec.template.spec.imagePullSecrets[]

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.initContainers

### Description

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

### Type

`array`

## .spec.template.spec.initContainers[]

### Description

A single application container that you want to run within a pod.

### Type

`object`

### Required

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |

| Property | Type | Description |
| --- | --- | --- |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.initContainers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].args[]

**Type**

`string`

## .spec.template.spec.initContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].command[]

**Type**

`string`

## .spec.template.spec.initContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

# .spec.template.spec.initContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

# .spec.template.spec.initContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

# .spec.template.spec.initContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.initContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.initContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.initContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.initContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.initContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.initContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.initContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.initContainers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.initContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.initContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.initContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.initContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.initContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.initContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.initContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.initContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.initContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.initContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa ↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.initContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.initContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.initContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.nodeSelector

**Description**

NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

**Type**

`object`

## .spec.template.spec.os

**Description**

PodOS defines the OS parameters of a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration ↗ Clients should expect to handle additional values and treat unrecognized values in this field as os: null |

## .spec.template.spec.overhead

**Description**

Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md

**Type**

`object`

## .spec.template.spec.readinessGates

**Description**

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates

**Type**

`array`

## .spec.template.spec.readinessGates[]

**Description**

PodReadinessGate contains the reference to a pod condition

**Type**

`object`

**Required**

`conditionType`

| Property | Type | Description |
| --- | --- | --- |
| `conditionType` | `string` | ConditionType refers to a condition in the pod's condition list with matching type. |

## .spec.template.spec.resourceClaims

**Description**

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable.

**Type**

`array`

## .spec.template.spec.resourceClaims[]

**Description**

PodResourceClaim references exactly one ResourceClaim, either directly or by naming a ResourceClaimTemplate which is then turned into a ResourceClaim for the pod. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL. |
| `resourceClaimName` | `string` | ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |
| `resourceClaimTemplateName` | `string` | ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.<br><br>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The pod name and resource name, along with a generated component, will be used to form a unique name for the ResourceClaim, which will be recorded in pod.status.resourceClaimStatuses.<br><br>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |

## .spec.template.spec.resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.template.spec.resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |

| Property | Type | Description |
| --- | --- | --- |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.schedulingGates

**Description**

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod. SchedulingGates can only be set at pod creation time, and be removed only afterwards.

**Type**

`array`

## .spec.template.spec.schedulingGates[]

**Description**

PodSchedulingGate is associated to a Pod to guard its scheduling.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the scheduling gate. Each scheduling gate must have a unique name field. |

## .spec.template.spec.securityContext

**Description**

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `fsGroup` | `integer` | A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:<br>1. The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw----<br>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows. |
| `fsGroupChangePolicy` | `string` | fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership(and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.<br>Possible enum values:<br>- `"Always"` indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior.<br>- `"OnRootMismatch"` indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `seLinuxChangePolicy` | `string` | seLinuxChangePolicy defines how the container's SELinux label is applied to all volumes used by the Pod. It has no effect on nodes that do not support SELinux or to volumes does not support SELinux. Valid values are "MountOption" and "Recursive".<br><br>"Recursive" means relabeling of all files on all Pod volumes by the container runtime. This may be slow for large volumes, but allows mixing privileged and unprivileged Pods sharing the same volume on the same node.<br><br>"MountOption" mounts all eligible Pod volumes with `-o context` mount option. This requires all Pods that share the same volume to use the same SELinux label. It is not possible to share the same volume among privileged and unprivileged Pods. Eligible volumes are in-tree FibreChannel and iSCSI volumes, and all CSI volumes whose CSI driver announces SELinux support by setting spec.seLinuxMount: true in their CSIDriver instance. Other volumes are always re-labelled recursively. "MountOption" value is allowed only when SELinuxMount feature gate is enabled.<br><br>If not specified and SELinuxMount feature gate is enabled, "MountOption" is used. If not specified and SELinuxMount feature gate is disabled, "MountOption" is used for ReadWriteOncePod volumes and "Recursive" for all other volumes.<br><br>This field affects only Pods that have SELinux label set, either in PodSecurityContext or in SecurityContext of all containers.<br><br>All Pods that use the same volume should use the same seLinuxChangePolicy, otherwise some pods can get stuck in ContainerCreating state. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `supplementalGroups` | `array` | A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows. |
| `supplementalGroupsPolicy` | `string` | Defines how supplemental groups of the first container processes are calculated. Valid values are "Merge" and "Strict". If not specified, "Merge" is used. (Alpha) Using the field requires the SupplementalGroupsPolicy feature gate to be enabled and the container runtime must implement support for this feature. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br><br>- `"Merge"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be merged with the primary user's groups as defined in the container image (in /etc/group). |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Strict"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be used instead of any groups defined in the container image. |
| `sysctls` | `array` | Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br>• `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>• `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `level` | `string` | Level is SELinux level label that applies to the container. |

| Property | Type | Description |
| --- | --- | --- |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.securityContext.supplementalGroups

**Description**

A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.supplementalGroups[]

**Type**

`integer`

## .spec.template.spec.securityContext.sysctls

**Description**

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.sysctls[]

**Description**

Sysctl defines a kernel parameter to be set

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of a property to set |
| `value` | `string` | Value of a property to set |

## .spec.template.spec.securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |

| Property | Type | Description |
| --- | --- | --- |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.tolerations

**Description**

If specified, the pod's tolerations.

**Type**

`array`

## .spec.template.spec.tolerations[]

**Description**

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `effect` | `string` | Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.<br><br>Possible enum values:<br>- `"NoExecute"` Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController.<br>- `"NoSchedule"` Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler.<br>- `"PreferNoSchedule"` Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler. |
| `key` | `string` | Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys. |
| `operator` | `string` | Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"Equal"`<br>- `"Exists"` |
| `tolerationSeconds` | `integer` | TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system. |
| `value` | `string` | Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string. |

## .spec.template.spec.topologySpreadConstraints

**Description**

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[]

**Description**

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

**Type**

`object`

**Required**

`maxSkew` `topologyKey` `whenUnsatisfiable`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.<br><br>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `maxSkew` | `integer` | MaxSkew describes the degree to which pods may be unevenly distributed. When `whenUnsatisfiable=DoNotSchedule`, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1. \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P \| - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When `whenUnsatisfiable=ScheduleAnyway`, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed. |
| `minDomains` | `integer` | MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.<br><br>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2: \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P P \| The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew. |
| `nodeAffinityPolicy` | `string` | NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.<br><br>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br><br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |
| `nodeTaintsPolicy` | `string` | NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.<br><br>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br><br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |

| Property | Type | Description |
| --- | --- | --- |
| `topologyKey` | `string` | TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field. |
| `whenUnsatisfiable` | `string` | WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: \| zone1 \| zone2 \| zone3 \| \| P P P \| P \| P \| If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it *more* imbalanced. It's a required field.<br><br>Possible enum values:<br>- `"DoNotSchedule"` instructs the scheduler not to schedule the pod when constraints are not satisfied.<br>- `"ScheduleAnyway"` instructs the scheduler to schedule the pod even if constraints are not satisfied. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector. This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.volumes

**Description**

List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

**Type**

`array`

## .spec.template.spec.volumes[]

**Description**

Volume represents a named volume in a pod that may be accessed by any container in the pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `awsElasticBlockStore` | `object` | Represents a Persistent Disk resource in AWS.<br><br>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling. |
| `azureDisk` | `object` | AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod. |
| `azureFile` | `object` | AzureFile represents an Azure File Service mount on the host and bind mount to the pod. |
| `cephfs` | `object` | Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling. |
| `cinder` | `object` | Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership |

| Property | Type | Description |
| --- | --- | --- |
| | | management and SELinux relabeling. |
| `configMap` | `object` | Adapts a ConfigMap into a volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling. |
| `csi` | `object` | Represents a source location of a volume to mount, managed by an external CSI driver |
| `downwardAPI` | `object` | DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling. |
| `emptyDir` | `object` | Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling. |
| `ephemeral` | `object` | Represents an ephemeral volume that is handled by a normal storage driver. |
| `fc` | `object` | Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling. |
| `flexVolume` | `object` | FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin. |
| `flocker` | `object` | Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling. |
| `gcePersistentDisk` | `object` | Represents a Persistent Disk resource in Google Compute Engine.<br><br>A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling. |
| `gitRepo` | `object` | Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.<br><br>DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container. |

| Property | Type | Description |
| --- | --- | --- |
| `glusterfs` | `object` | Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling. |
| `hostPath` | `object` | Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling. |
| `image` | `object` | ImageVolumeSource represents a image volume resource. |
| `iscsi` | `object` | Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling. |
| `name` | `string` | name of the volume. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `nfs` | `object` | Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling. |
| `persistentVolumeClaim` | `object` | PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system). |
| `photonPersistentDisk` | `object` | Represents a Photon Controller persistent disk resource. |
| `portworxVolume` | `object` | PortworxVolumeSource represents a Portworx volume resource. |
| `projected` | `object` | Represents a projected volume source |
| `quobyte` | `object` | Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling. |
| `rbd` | `object` | Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling. |
| `scaleIO` | `object` | ScaleIOVolumeSource represents a persistent ScaleIO volume |

| Property | Type | Description |
| --- | --- | --- |
| `secret` | `object` | Adapts a Secret into a volume.<br><br>The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling. |
| `storageos` | `object` | Represents a StorageOS persistent volume resource. |
| `vsphereVolume` | `object` | Represents a vSphere volume resource. |

## .spec.template.spec.volumes[].awsElasticBlockStore

**Description**

Represents a Persistent Disk resource in AWS. An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). |
| `readOnly` | `boolean` | readOnly value true will force the readOnly setting in VolumeMounts. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `volumeID` | `string` | volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |

## .spec.template.spec.volumes[].azureDisk

**Description**

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`diskName` `diskURI`

| Property | Type | Description |
| --- | --- | --- |
| `cachingMode` | `string` | cachingMode is the Host Caching mode: None, Read Only, Read Write.<br><br>Possible enum values:<br>• `"None"`<br>• `"ReadOnly"`<br>• `"ReadWrite"` |
| `diskName` | `string` | diskName is the Name of the data disk in the blob storage |
| `diskURI` | `string` | diskURI is the URI of data disk in the blob storage |
| `fsType` | `string` | fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `kind` | `string` | kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared<br><br>Possible enum values:<br>• `"Dedicated"`<br>• `"Managed"`<br>• `"Shared"` |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

## .spec.template.spec.volumes[].azureFile

**Description**

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`secretName` `shareName`

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretName` | `string` | secretName is the name of secret that contains Azure Storage Account Name and Key |
| `shareName` | `string` | shareName is the azure share Name |

## .spec.template.spec.volumes[].cephfs

**Description**

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`monitors`

| Property | Type | Description |
| --- | --- | --- |
| `monitors` | `array` | monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `path` | `string` | path is Optional: Used as the mounted root, rather than the full Ceph tree, default is / |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretFile` | `string` | secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].cephfs.monitors

**Description**

monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].cephfs.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].cephfs.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].cinder

**Description**

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeID` | `string` | volumeID used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |

## .spec.template.spec.volumes[].cinder.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].configMap

**Description**

Adapts a ConfigMap into a volume. The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `defaultMode` | `integer` | defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].csi

**Description**

Represents a source location of a volume to mount, managed by an external CSI driver

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
|---|---|---|
| `driver` | `string` | driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster. |
| `fsType` | `string` | fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply. |

| Property | Type | Description |
| --- | --- | --- |
| `nodePublishSecretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `readOnly` | `boolean` | readOnly specifies a read-only configuration for the volume. Defaults to false (read/write). |
| `volumeAttributes` | `object` | volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values. |

## .spec.template.spec.volumes[].csi.nodePublishSecretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].csi.volumeAttributes

**Description**

volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**Type**

`object`

## .spec.template.spec.volumes[].downwardAPI

**Description**

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | Items is a list of downward API volume file |

## .spec.template.spec.volumes[].downwardAPI.items

**Description**

Items is a list of downward API volume file

**Type**

`array`

## .spec.template.spec.volumes[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
|---|---|---|
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].emptyDir

**Description**

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `medium` | `string` | medium represents what type of storage medium should back this directory. The default is "" which means to use the n |
| `sizeLimit` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAM<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits> <r<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.html)<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a number<br>When a Quantity is parsed from a string, it will remember the type of suffix it had, an<br>Before serializing, Quantity will be put in "canonical form". This means that Exponent/<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suffix<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point number.<br>Non-canonical values will still parse as long as they are well formed, but will be re-e<br>This format is intended to make it difficult to use these numbers without writing some |

## .spec.template.spec.volumes[].ephemeral

**Description**

Represents an ephemeral volume that is handled by a normal storage driver.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `volumeClaimTemplate` | `object` | PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate

**Description**

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

**Type**

`object`

**Required**

`spec`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec

### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `accessModes` | `array` | accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1 ↗ |
| `dataSource` | `object` | TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace. |
| `dataSourceRef` | `object` | TypedObjectReference contains enough information to let you locate the typed referenced object |
| `resources` | `object` | VolumeResourceRequirements describes the storage resource requirements for a volume. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `storageClassName` | `string` | storageClassName is the name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1 ↗ |
| `volumeAttributesClassName` | `string` | volumeAttributesClassName may be used to set the VolumeAttributesClass used by this claim. If specified, the CSI driver will create or update the volume with the attributes defined in the corresponding VolumeAttributesClass. This has a different purpose than storageClassName, it can be changed after the claim is created. An empty string value means that no VolumeAttributesClass will be applied to the claim but it's not allowed to reset this field to empty string once it is set. If unspecified and the PersistentVolumeClaim is unbound, the default VolumeAttributesClass will be set by the persistentvolume controller if it exists. If the resource referred to by volumeAttributesClass does not exist, this PersistentVolumeClaim will be set to a Pending state, as reflected by the |

| Property | Type | Description |
| --- | --- | --- |
| | | modifyVolumeStatus field, until such as a resource exists. More info: https://kubernetes.io/docs/concepts/storage/volume-attributes-classes/ (Beta) Using this field requires the VolumeAttributesClass feature gate to be enabled (off by default). |
| `volumeMode` | `string` | volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.<br>Possible enum values:<br>• `"Block"` means the volume will not be formatted with a filesystem and will remain a raw block device.<br>• `"Filesystem"` means the volume will be or is formatted with a filesystem. |
| `volumeName` | `string` | volumeName is the binding reference to the PersistentVolume backing this claim. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes

**Description**

accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

**Description**

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is the name of resource being referenced |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef

**Description**

TypedObjectReference contains enough information to let you locate the typed referenced object

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |
| `namespace` | `string` | Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

**Description**

VolumeResourceRequirements describes the storage resource requirements for a volume.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].fc

**Description**

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `lun` | `integer` | lun is Optional: FC target lun number |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `targetWWNs` | `array` | targetWWNs is Optional: FC target worldwide names (WWNs) |
| `wwids` | `array` | wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously. |

## .spec.template.spec.volumes[].fc.targetWWNs

**Description**

targetWWNs is Optional: FC target worldwide names (WWNs)

**Type**

`array`

## .spec.template.spec.volumes[].fc.targetWWNs[]

**Type**

`string`

## .spec.template.spec.volumes[].fc.wwids

**Description**

wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

**Type**

`array`

## .spec.template.spec.volumes[].fc.wwids[]

**Type**

`string`

## .spec.template.spec.volumes[].flexVolume

**Description**

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
| --- | --- | --- |
| `driver` | `string` | driver is the name of the driver to use for this volume. |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script. |
| `options` | `object` | options is Optional: this field holds extra command options if any. |
| `readOnly` | `boolean` | readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |

## .spec.template.spec.volumes[].flexVolume.options

**Description**

options is Optional: this field holds extra command options if any.

**Type**

`object`

## .spec.template.spec.volumes[].flexVolume.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: |

| Property | Type | Description |
| --- | --- | --- |
| | | https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].flocker

### Description

Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `datasetName` | `string` | datasetName is Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated |
| `datasetUUID` | `string` | datasetUUID is the UUID of the dataset. This is unique identifier of a Flocker dataset |

## .spec.template.spec.volumes[].gcePersistentDisk

### Description

Represents a Persistent Disk resource in Google Compute Engine. A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

### Type

`object`

### Required

`pdName`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `pdName` | `string` | pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

## .spec.template.spec.volumes[].gitRepo

**Description**

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

**Type**

`object`

**Required**

`repository`

| Property | Type | Description |
| --- | --- | --- |
| `directory` | `string` | directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name. |
| `repository` | `string` | repository is the URL |
| `revision` | `string` | revision is the commit hash for the specified revision. |

## .spec.template.spec.volumes[].glusterfs

**Description**

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`endpoints` `path`

| Property | Type | Description |
| --- | --- | --- |
| `endpoints` | `string` | endpoints is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |
| `path` | `string` | path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

## .spec.template.spec.volumes[].hostPath

### Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`path`

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗ |
| `type` | `string` | type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗<br><br>Possible enum values:<br>• `""` For backwards compatible, leave it empty if unset<br>• `"BlockDevice"` A block device must exist at the given path<br>• `"CharDevice"` A character device must exist at the given path<br>• `"Directory"` A directory must exist at the given path<br>• `"DirectoryOrCreate"` If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet.<br>• `"File"` A file must exist at the given path<br>• `"FileOrCreate"` If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet.<br>• `"Socket"` A UNIX socket must exist at the given path |

## .spec.template.spec.volumes[].image

### Description

ImageVolumeSource represents a image volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `pullPolicy` | `string` | Policy for pulling OCI objects. Possible values are: Always: the kubelet always attempts to pull the reference. Container creation will fail If the pull fails. Never: the kubelet never pulls the reference and only uses a local image or artifact. |

| Property | Type | Description |
| --- | --- | --- |
| | | Container creation will fail if the reference isn't present. IfNotPresent: the kubelet pulls if the reference isn't already present on disk. Container creation will fail if the reference isn't present and the pull fails. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise.<br><br>Possible enum values:<br><br>- `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>- `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>- `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `reference` | `string` | Required: Image or artifact reference to be used. Behaves in the same way as pod.spec.containers[*].image. Pull secrets will be assembled in the same way as for the container image by looking up node credentials, SA image pull secrets, and pod spec image pull secrets. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |

## .spec.template.spec.volumes[].iscsi

### Description

Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`targetPortal` `iqn` `lun`

| Property | Type | Description |
| --- | --- | --- |
| `chapAuthDiscovery` | `boolean` | chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication |
| `chapAuthSession` | `boolean` | chapAuthSession defines whether support iSCSI Session CHAP authentication |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi |
| `initiatorName` | `string` | initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection. |
| `iqn` | `string` | iqn is the target iSCSI Qualified Name. |

| Property | Type | Description |
| --- | --- | --- |
| `iscsiInterface` | `string` | iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp). |
| `lun` | `integer` | lun represents iSCSI Target Lun number. |
| `portals` | `array` | portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `targetPortal` | `string` | targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |

## .spec.template.spec.volumes[].iscsi.portals

**Description**

portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

**Type**

`array`

## .spec.template.spec.volumes[].iscsi.portals[]

**Type**

`string`

## .spec.template.spec.volumes[].iscsi.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].nfs

**Description**

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`server` `path`

| Property | Type | Description |
|---|---|---|
| `path` | `string` | path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |
| `readOnly` | `boolean` | readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |
| `server` | `string` | server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |

## .spec.template.spec.volumes[].persistentVolumeClaim

**Description**

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

**Type**

`object`

**Required**

`claimName`

| Property | Type | Description |
|---|---|---|
| `claimName` | `string` | claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims |
| `readOnly` | `boolean` | readOnly Will force the ReadOnly setting in VolumeMounts. Default false. |

## .spec.template.spec.volumes[].photonPersistentDisk

**Description**

Represents a Photon Controller persistent disk resource.

**Type**

`object`

**Required**

`pdID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `pdID` | `string` | pdID is the ID that identifies Photon Controller persistent disk |

## .spec.template.spec.volumes[].portworxVolume

**Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `volumeID` | `string` | volumeID uniquely identifies a Portworx volume |

## .spec.template.spec.volumes[].projected

**Description**

Represents a projected volume source

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `sources` | `array` | sources is the list of volume projections. Each entry in this list handles one source. |

## .spec.template.spec.volumes[].projected.sources

**Description**

sources is the list of volume projections. Each entry in this list handles one source.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[]

**Description**

Projection that may be projected along with other supported volume types. Exactly one of these fields must be set.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `clusterTrustBundle` | `object` | ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem. |
| `configMap` | `object` | Adapts a ConfigMap into a projected volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode. |
| `downwardAPI` | `object` | Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode. |
| `secret` | `object` | Adapts a secret into a projected volume.<br><br>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode. |
| `serviceAccountToken` | `object` | ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise). |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle

**Description**

ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem.

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `name` | `string` | Select a single ClusterTrustBundle by object name. Mutually-exclusive with signerName and labelSelector. |
| `optional` | `boolean` | If true, don't block pod startup if the referenced ClusterTrustBundle(s) aren't available. If using name, then the named ClusterTrustBundle is allowed not to exist. If using signerName, then the combination of signerName and labelSelector is allowed to match zero ClusterTrustBundles. |
| `path` | `string` | Relative path from the volume root to write the bundle. |
| `signerName` | `string` | Select all ClusterTrustBundles that match this signer name. Mutually-exclusive with name. The contents of all selected ClusterTrustBundles will be unified and deduplicated. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].projected.sources[].configMap

**Description**

Adapts a ConfigMap into a projected volume. The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].projected.sources[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI

**Description**

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | Items is a list of DownwardAPIVolume file |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items

**Description**

Items is a list of DownwardAPIVolume file

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |

| Property | Type | Description |
| --- | --- | --- |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and |

| Property | Type | Description |
| --- | --- | --- |
| | | The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].projected.sources[].secret

### Description

Adapts a secret into a projected volume. The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | optional field specify whether the Secret or its key must be defined |

## .spec.template.spec.volumes[].projected.sources[].secret.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].serviceAccountToken

**Description**

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `audience` | `string` | audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver. |
| `expirationSeconds` | `integer` | expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours.Defaults to 1 hour and must be at least 10 minutes. |
| `path` | `string` | path is the path relative to the mount point of the file to project the token into. |

## .spec.template.spec.volumes[].quobyte

**Description**

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`registry` `volume`

| Property | Type | Description |
| --- | --- | --- |
| `group` | `string` | group to map volume access to Default is no group |
| `readOnly` | `boolean` | readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false. |
| `registry` | `string` | registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes |
| `tenant` | `string` | tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin |
| `user` | `string` | user to map volume access to Defaults to serivceaccount user |

| Property | Type | Description |
| --- | --- | --- |
| `volume` | `string` | volume is a string that references an already created Quobyte volume by name. |

## .spec.template.spec.volumes[].rbd

### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`monitors` `image`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd ↗ |
| `image` | `string` | image is the rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `keyring` | `string` | keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `monitors` | `array` | monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `pool` | `string` | pool is the rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is the rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].rbd.monitors

**Description**

monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].rbd.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].rbd.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].scaleIO

**Description**

ScaleIOVolumeSource represents a persistent ScaleIO volume

**Type**

`object`

**Required**

`gateway` `system` `secretRef`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs". |
| `gateway` | `string` | gateway is the host address of the ScaleIO API Gateway. |
| `protectionDomain` | `string` | protectionDomain is the name of the ScaleIO Protection Domain for the configured storage. |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

| Property | Type | Description |
| --- | --- | --- |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `sslEnabled` | `boolean` | sslEnabled Flag enable/disable SSL communication with Gateway, default false |
| `storageMode` | `string` | storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned. |
| `storagePool` | `string` | storagePool is the ScaleIO Storage Pool associated with the protection domain. |
| `system` | `string` | system is the name of the storage system as configured in ScaleIO. |
| `volumeName` | `string` | volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source. |

## .spec.template.spec.volumes[].scaleIO.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].secret

**Description**

Adapts a Secret into a volume. The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON |

| Property | Type | Description |
| --- | --- | --- |
| | | requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `optional` | `boolean` | optional field specify whether the Secret or its keys must be defined |
| `secretName` | `string` | secretName is the name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret ↗ |

## .spec.template.spec.volumes[].secret.items

**Description**

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].storageos

### Description

Represents a StorageOS persistent volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeName` | `string` | volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace. |
| `volumeNamespace` | `string` | volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created. |

## .spec.template.spec.volumes[].storageos.secretRef

### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].vsphereVolume

**Description**

Represents a vSphere volume resource.

**Type**

`object`

**Required**

`volumePath`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `storagePolicyID` | `string` | storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName. |
| `storagePolicyName` | `string` | storagePolicyName is the storage Policy Based Management (SPBM) profile name. |
| `volumePath` | `string` | volumePath is the path that identifies vSphere volume vmdk |

## .spec.updateStrategy

**Description**

DaemonSetUpdateStrategy is a struct used to control the update strategy for a DaemonSet.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `rollingUpdate` | `object` | Spec to control the desired behavior of daemon set rolling update. |
| `type` | `string` | Type of daemon set update. Can be "RollingUpdate" or "OnDelete". Default is RollingUpdate.<br><br>Possible enum values:<br>- `"OnDelete"` Replace the old daemons only when it's killed<br>- `"RollingUpdate"` Replace the old daemons by new ones using rolling update i.e replace them on each node one after the other. |

## .spec.updateStrategy.rollingUpdate

**Description**

Spec to control the desired behavior of daemon set rolling update.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `maxSurge` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `maxUnavailable` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .status

**Description**

DaemonSetStatus represents the current status of a daemon set.

**Type**

`object`

**Required**

`currentNumberScheduled` `numberMisscheduled` `desiredNumberScheduled` `numberReady`

| Property | Type | Description |
| --- | --- | --- |
| `collisionCount` | `integer` | Count of hash collisions for the DaemonSet. The DaemonSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision. |
| `conditions` | `array` | Represents the latest available observations of a DaemonSet's current state. |
| `currentNumberScheduled` | `integer` | The number of nodes that are running at least 1 daemon pod and are supposed to run the daemon pod. More info: https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/ |
| `desiredNumberScheduled` | `integer` | The total number of nodes that should be running the daemon pod (including nodes correctly running the daemon pod). More info: https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/ |
| `numberAvailable` | `integer` | The number of nodes that should be running the daemon pod and have one or more of the daemon pod running and available (ready for at least spec.minReadySeconds) |
| `numberMisscheduled` | `integer` | The number of nodes that are running the daemon pod, but are not supposed to run the daemon pod. More info: https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/ |

| Property | Type | Description |
| --- | --- | --- |
| `numberReady` | `integer` | numberReady is the number of nodes that should be running the daemon pod and have one or more of the daemon pod running with a Ready Condition. |
| `numberUnavailable` | `integer` | The number of nodes that should be running the daemon pod and have none of the daemon pod running and available (ready for at least spec.minReadySeconds) |
| `observedGeneration` | `integer` | The most recent generation observed by the daemon set controller. |
| `updatedNumberScheduled` | `integer` | The total number of nodes that are running updated daemon pod |

## .status.conditions

**Description**

Represents the latest available observations of a DaemonSet's current state.

**Type**

`array`

## .status.conditions[]

**Description**

DaemonSetCondition describes the state of a DaemonSet at a certain point.

**Type**

`object`

**Required**

`type` `status`

| Property | Type | Description |
| --- | --- | --- |
| `lastTransitionTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | A human readable message indicating details about the transition. |
| `reason` | `string` | The reason for the condition's last transition. |
| `status` | `string` | Status of the condition, one of True, False, Unknown. |

| Property | Type | Description |
| --- | --- | --- |
| `type` | `string` | Type of DaemonSet condition. |

# API Endpoints

The following API endpoints are available:

- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/daemonsets`
  - `DELETE` : delete collection of DaemonSet
  - `GET` : list objects of kind DaemonSet
  - `POST` : create a new DaemonSet
- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/daemonsets/{name}`
  - `DELETE` : delete the specified DaemonSet
  - `GET` : read the specified DaemonSet
  - `PATCH` : partially update the specified DaemonSet
  - `PUT` : replace the specified DaemonSet
- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/daemonsets/{name}/status`
  - `GET` : read status of the specified DaemonSet
  - `PATCH` : partially update status of the specified DaemonSet
  - `PUT` : replace status of the specified DaemonSet

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/daemonsets

**HTTP method**

`DELETE`

**Description**

delete collection of DaemonSet

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`GET`

**Description**

list objects of kind DaemonSet

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `DaemonSetList` schema |

| HTTP code | Response body |
| --- | --- |
| 401 - Unauthorized | Empty |

### HTTP method

`POST`

### Description

create a new DaemonSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `DaemonSet` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `DaemonSet` schema |
| 201 - Created | `DaemonSet` schema |
| 202 - Accepted | `DaemonSet` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/daemonsets/{name}

### HTTP method

`DELETE`

### Description

delete the specified DaemonSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 202 - Accepted | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

### Description

read the specified DaemonSet

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `DaemonSet` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update the specified DaemonSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `DaemonSet` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace the specified DaemonSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `DaemonSet` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `DaemonSet` schema |
| 201 - Created | `DaemonSet` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/daemonsets/{name}/status

### HTTP method

`GET`

### Description

read status of the specified DaemonSet

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `DaemonSet` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update status of the specified DaemonSet

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `DaemonSet` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace status of the specified DaemonSet

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
|---|---|---|
| `body` | `DaemonSet` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `DaemonSet` schema |
| 201 - Created | `DaemonSet` schema |
| 401 - Unauthorized | Empty |

# Deployment [apps/v1]

**Description**

Deployment enables declarative updates for Pods and ReplicaSets.

**Type**

`object`

## Specification

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources |
| `kind` | `string` | Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | DeploymentSpec is the specification of the desired behavior of the Deployment. |
| `status` | `object` | DeploymentStatus is the most recently observed status of the Deployment. |

### .spec

**Description**

DeploymentSpec is the specification of the desired behavior of the Deployment.

**Type**

`object`

**Required**

`selector` `template`

| Property | Type | Description |
| --- | --- | --- |
| `minReadySeconds` | `integer` | Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready) |
| `paused` | `boolean` | Indicates that the deployment is paused. |
| `progressDeadlineSeconds` | `integer` | The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s. |
| `replicas` | `integer` | Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1. |
| `revisionHistoryLimit` | `integer` | The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `strategy` | `object` | DeploymentStrategy describes how to replace existing pods with new ones. |
| `template` | `object` | PodTemplateSpec describes the data a pod should have when created from a template |

## .spec.selector

### Description

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only |

| Property | Type | Description |
| --- | --- | --- |
| | | "value". The requirements are ANDed. |

## .spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.strategy

**Description**

DeploymentStrategy describes how to replace existing pods with new ones.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `rollingUpdate` | `object` | Spec to control the desired behavior of rolling update. |
| `type` | `string` | Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.<br><br>Possible enum values:<br>- `"Recreate"` Kill all existing pods before creating new ones.<br>- `"RollingUpdate"` Replace the old ReplicaSets by new one using rolling update i.e gradually scale down the old ReplicaSets and scale up the new one. |

## .spec.strategy.rollingUpdate

**Description**

Spec to control the desired behavior of rolling update.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `maxSurge` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `maxUnavailable` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template

**Description**

PodTemplateSpec describes the data a pod should have when created from a template

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PodSpec is a description of a pod. |

## .spec.template.spec

### Description

PodSpec is a description of a pod.

### Type

`object`

### Required

`containers`

| Property | Type | Description |
| --- | --- | --- |
| `activeDeadlineSeconds` | `integer` | Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer. |
| `affinity` | `object` | Affinity is a group of affinity scheduling rules. |
| `automountServiceAccountToken` | `boolean` | AutomountServiceAccountToken indicates whether a service account token should be automatically mounted. |
| `containers` | `array` | List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated. |
| `dnsConfig` | `object` | PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy. |
| `dnsPolicy` | `string` | Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.<br><br>Possible enum values:<br><br>- `"ClusterFirst"` indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"ClusterFirstWithHostNet"` indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>• `"Default"` indicates that the pod should use the default (as determined by kubelet) DNS settings.<br>• `"None"` indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig. |
| `enableServiceLinks` | `boolean` | EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching traditional container linking syntax. Optional: Defaults to true. |
| `ephemeralContainers` | `array` | List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. |
| `hostAliases` | `array` | HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. |
| `hostIPC` | `boolean` | Use the host's ipc namespace. Optional: Default to false. |
| `hostNetwork` | `boolean` | Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false. |
| `hostPID` | `boolean` | Use the host's pid namespace. Optional: Default to false. |
| `hostUsers` | `boolean` | Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new userns is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature. |
| `hostname` | `string` | Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value. |
| `imagePullSecrets` | `array` | ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `initContainers` | `array` | List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/ |
| `nodeName` | `string` | NodeName indicates in which node this pod is scheduled. If empty, this pod is a candidate for scheduling by the scheduler defined in schedulerName. Once this field is set, the kubelet for this node becomes responsible for the lifecycle of this pod. This field should not be used to express a desire for the pod to be scheduled on a specific node. https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodename |
| `nodeSelector` | `object` | NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/ |
| `os` | `object` | PodOS defines the OS parameters of a pod. |
| `overhead` | `object` | Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md |
| `preemptionPolicy` | `string` | PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.<br><br>Possible enum values:<br>• `"Never"` means that pod never preempts other pods with lower priority.<br>• `"PreemptLowerPriority"` means that pod can preempt other pods with lower priority. |
| `priority` | `integer` | The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority. |

| Property | Type | Description |
| --- | --- | --- |
| `priorityClassName` | `string` | If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default. |
| `readinessGates` | `array` | If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates ↗ |
| `resourceClaims` | `array` | ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.<br><br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy ↗<br><br>Possible enum values:<br>• `"Always"`<br>• `"Never"`<br>• `"OnFailure"` |
| `runtimeClassName` | `string` | RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class ↗ |
| `schedulerName` | `string` | If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler. |
| `schedulingGates` | `array` | SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.<br><br>SchedulingGates can only be set at pod creation time, and be removed only afterwards. |

| Property | Type | Description |
| --- | --- | --- |
| `securityContext` | `object` | PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext. |
| `serviceAccount` | `string` | DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead. |
| `serviceAccountName` | `string` | ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/ ↗ |
| `setHostnameAsFQDN` | `boolean` | If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false. |
| `shareProcessNamespace` | `boolean` | Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false. |
| `subdomain` | `string` | If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all. |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds. |
| `tolerations` | `array` | If specified, the pod's tolerations. |
| `topologySpreadConstraints` | `array` | TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed. |

| Property | Type | Description |
| --- | --- | --- |
| `volumes` | `array` | List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes ↗ |

## .spec.template.spec.affinity

**Description**

Affinity is a group of affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `nodeAffinity` | `object` | Node affinity is a group of node affinity scheduling rules. |
| `podAffinity` | `object` | Pod affinity is a group of inter pod affinity scheduling rules. |
| `podAntiAffinity` | `object` | Pod anti affinity is a group of inter pod anti affinity scheduling rules. |

## .spec.template.spec.affinity.nodeAffinity

**Description**

Node affinity is a group of node affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `object` | A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

**Type**

`object`

**Required**

`weight` `preference`

| Property | Type | Description |
|---|---|---|
| `preference` | `object` | A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm. |
| `weight` | `integer` | Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

**Type**

`object`

**Required**

`nodeSelectorTerms`

| Property | Type | Description |
| --- | --- | --- |
| `nodeSelectorTerms` | `array` | Required. A list of node selector terms. The terms are ORed. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms

**Description**

Required. A list of node selector terms. The terms are ORed.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |

| Property | Type | Description |
| --- | --- | --- |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity

**Description**

Pod affinity is a group of inter pod affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the |

| Property | Type | Description |
| --- | --- | --- |
| | | corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
| --- | --- | --- |
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into |

| Property | Type | Description |
| --- | --- | --- |
| | | consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity

**Description**

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the |

| Property | Type | Description |
| --- | --- | --- |
| | | elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
| --- | --- | --- |
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

# .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into |

| Property | Type | Description |
| --- | --- | --- |
| | | consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.containers

**Description**

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[]

**Description**

A single application container that you want to run within a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |

| Property | Type | Description |
| --- | --- | --- |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |

| Property | Type | Description |
| --- | --- | --- |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.containers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].args[]

**Type**

`string`

## .spec.template.spec.containers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].command[]

**Type**

`string`

## .spec.template.spec.containers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.containers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.containers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.containers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.containers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.containers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.containers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.containers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.containers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.containers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.containers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.containers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.containers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.containers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.containers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.containers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.containers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.containers[].securityContext.appArmorProfile

### Description

AppArmorProfile defines a pod or container's AppArmor settings.

### Type

`object`

### Required

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.containers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.containers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.containers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.containers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.containers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.containers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.containers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.containers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.dnsConfig

**Description**

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `nameservers` | `array` | A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed. |
| `options` | `array` | A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy. |
| `searches` | `array` | A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed. |

## .spec.template.spec.dnsConfig.nameservers

**Description**

A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.nameservers[]

**Type**

`string`

## .spec.template.spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

`array`

## .spec.template.spec.dnsConfig.options[]

**Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is this DNS resolver option's name. Required. |
| `value` | `string` | Value is this DNS resolver option's value. |

## .spec.template.spec.dnsConfig.searches

**Description**

A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.searches[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers

**Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[]

**Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable |

| Property | Type | Description |
| --- | --- | --- |
| | | exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | Ports are not allowed for ephemeral containers. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for the container to manage the restart behavior of each container within a pod. This may only be set for init containers. You cannot set this field on ephemeral containers. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `targetContainerName` | `string` | If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.<br><br>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined. |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message |

| Property | Type | Description |
| --- | --- | --- |
| | | length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>- `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>- `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.ephemeralContainers[].args

**Description**

Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].args[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.ephemeralContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.ephemeralContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

# .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

# .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

# .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].ports

**Description**

Ports are not allowed for ephemeral containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |

| Property | Type | Description |
| --- | --- | --- |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br><br>Possible enum values:<br>- `"SCTP"` is the SCTP protocol.<br>- `"TCP"` is the TCP protocol.<br>- `"UDP"` is the UDP protocol. |

## .spec.template.spec.ephemeralContainers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.ephemeralContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.ephemeralContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.ephemeralContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.ephemeralContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.ephemeralContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.ephemeralContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.ephemeralContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.ephemeralContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.hostAliases

**Description**

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified.

**Type**

`array`

## .spec.template.spec.hostAliases[]

**Description**

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

**Type**

`object`

**Required**

`ip`

| Property | Type | Description |
|---|---|---|
| `hostnames` | `array` | Hostnames for the above IP address. |
| `ip` | `string` | IP address of the host file entry. |

## .spec.template.spec.hostAliases[].hostnames

**Description**

Hostnames for the above IP address.

**Type**

`array`

## .spec.template.spec.hostAliases[].hostnames[]

**Type**

`string`

## .spec.template.spec.imagePullSecrets

**Description**

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

**Type**

`array`

## .spec.template.spec.imagePullSecrets[]

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.initContainers

### Description

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

### Type

`array`

## .spec.template.spec.initContainers[]

### Description

A single application container that you want to run within a pod.

### Type

`object`

### Required

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |

| Property | Type | Description |
| --- | --- | --- |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.initContainers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].args[]

**Type**

`string`

## .spec.template.spec.initContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].command[]

**Type**

`string`

## .spec.template.spec.initContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

# .spec.template.spec.initContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

# .spec.template.spec.initContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

# .spec.template.spec.initContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.initContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.initContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.initContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.initContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.initContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.initContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.initContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.initContainers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.initContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.initContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.initContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.initContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.initContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.initContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.initContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.initContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.initContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.initContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.initContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.initContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.initContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.nodeSelector

**Description**

NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

**Type**

`object`

## .spec.template.spec.os

**Description**

PodOS defines the OS parameters of a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration ↗ Clients should expect to handle additional values and treat unrecognized values in this field as os: null |

## .spec.template.spec.overhead

**Description**

Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md

**Type**

`object`

## .spec.template.spec.readinessGates

**Description**

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates

**Type**

`array`

## .spec.template.spec.readinessGates[]

**Description**

PodReadinessGate contains the reference to a pod condition

**Type**

`object`

**Required**

`conditionType`

| Property | Type | Description |
|---|---|---|
| `conditionType` | `string` | ConditionType refers to a condition in the pod's condition list with matching type. |

## .spec.template.spec.resourceClaims

**Description**

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable.

**Type**

`array`

## .spec.template.spec.resourceClaims[]

**Description**

PodResourceClaim references exactly one ResourceClaim, either directly or by naming a ResourceClaimTemplate which is then turned into a ResourceClaim for the pod. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL. |
| `resourceClaimName` | `string` | ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |
| `resourceClaimTemplateName` | `string` | ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.<br><br>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The pod name and resource name, along with a generated component, will be used to form a unique name for the ResourceClaim, which will be recorded in pod.status.resourceClaimStatuses.<br><br>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |

## .spec.template.spec.resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.template.spec.resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |

| Property | Type | Description |
| --- | --- | --- |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.schedulingGates

**Description**

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod. SchedulingGates can only be set at pod creation time, and be removed only afterwards.

**Type**

`array`

## .spec.template.spec.schedulingGates[]

**Description**

PodSchedulingGate is associated to a Pod to guard its scheduling.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the scheduling gate. Each scheduling gate must have a unique name field. |

## .spec.template.spec.securityContext

**Description**

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `fsGroup` | `integer` | A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:<br>1. The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw----<br>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows. |
| `fsGroupChangePolicy` | `string` | fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership(and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.<br>Possible enum values:<br>- `"Always"` indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior.<br>- `"OnRootMismatch"` indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `seLinuxChangePolicy` | `string` | seLinuxChangePolicy defines how the container's SELinux label is applied to all volumes used by the Pod. It has no effect on nodes that do not support SELinux or to volumes does not support SELinux. Valid values are "MountOption" and "Recursive".<br><br>"Recursive" means relabeling of all files on all Pod volumes by the container runtime. This may be slow for large volumes, but allows mixing privileged and unprivileged Pods sharing the same volume on the same node.<br><br>"MountOption" mounts all eligible Pod volumes with `-o context` mount option. This requires all Pods that share the same volume to use the same SELinux label. It is not possible to share the same volume among privileged and unprivileged Pods. Eligible volumes are in-tree FibreChannel and iSCSI volumes, and all CSI volumes whose CSI driver announces SELinux support by setting spec.seLinuxMount: true in their CSIDriver instance. Other volumes are always re-labelled recursively. "MountOption" value is allowed only when SELinuxMount feature gate is enabled.<br><br>If not specified and SELinuxMount feature gate is enabled, "MountOption" is used. If not specified and SELinuxMount feature gate is disabled, "MountOption" is used for ReadWriteOncePod volumes and "Recursive" for all other volumes.<br><br>This field affects only Pods that have SELinux label set, either in PodSecurityContext or in SecurityContext of all containers.<br><br>All Pods that use the same volume should use the same seLinuxChangePolicy, otherwise some pods can get stuck in ContainerCreating state. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `supplementalGroups` | `array` | A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows. |
| `supplementalGroupsPolicy` | `string` | Defines how supplemental groups of the first container processes are calculated. Valid values are "Merge" and "Strict". If not specified, "Merge" is used. (Alpha) Using the field requires the SupplementalGroupsPolicy feature gate to be enabled and the container runtime must implement support for this feature. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br><br>- `"Merge"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be merged with the primary user's groups as defined in the container image (in /etc/group). |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Strict"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be used instead of any groups defined in the container image. |
| `sysctls` | `array` | Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br>• `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>• `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `level` | `string` | Level is SELinux level label that applies to the container. |

| Property | Type | Description |
| --- | --- | --- |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.securityContext.supplementalGroups

**Description**

A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.supplementalGroups[]

**Type**

`integer`

## .spec.template.spec.securityContext.sysctls

**Description**

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.sysctls[]

**Description**

Sysctl defines a kernel parameter to be set

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of a property to set |
| `value` | `string` | Value of a property to set |

## .spec.template.spec.securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |

| Property | Type | Description |
| --- | --- | --- |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.tolerations

**Description**

If specified, the pod's tolerations.

**Type**

`array`

## .spec.template.spec.tolerations[]

**Description**

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `effect` | `string` | Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.<br><br>Possible enum values:<br>- `"NoExecute"` Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController.<br>- `"NoSchedule"` Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler.<br>- `"PreferNoSchedule"` Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler. |
| `key` | `string` | Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys. |
| `operator` | `string` | Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"Equal"`<br>- `"Exists"` |
| `tolerationSeconds` | `integer` | TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system. |
| `value` | `string` | Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string. |

## .spec.template.spec.topologySpreadConstraints

**Description**

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[]

**Description**

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

**Type**

`object`

**Required**

`maxSkew` `topologyKey` `whenUnsatisfiable`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.<br><br>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `maxSkew` | `integer` | MaxSkew describes the degree to which pods may be unevenly distributed. When `whenUnsatisfiable=DoNotSchedule`, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1. \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P \| - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When `whenUnsatisfiable=ScheduleAnyway`, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed. |
| `minDomains` | `integer` | MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.<br><br>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2: \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P P \| The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew. |
| `nodeAffinityPolicy` | `string` | NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.<br><br>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br><br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |
| `nodeTaintsPolicy` | `string` | NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.<br><br>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br><br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |

| Property | Type | Description |
| --- | --- | --- |
| `topologyKey` | `string` | TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field. |
| `whenUnsatisfiable` | `string` | WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: \| zone1 \| zone2 \| zone3 \| \| P P P \| P \| P \| If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it *more* imbalanced. It's a required field.<br><br>Possible enum values:<br>- `"DoNotSchedule"` instructs the scheduler not to schedule the pod when constraints are not satisfied.<br>- `"ScheduleAnyway"` instructs the scheduler to schedule the pod even if constraints are not satisfied. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector. This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.volumes

**Description**

List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

**Type**

`array`

## .spec.template.spec.volumes[]

**Description**

Volume represents a named volume in a pod that may be accessed by any container in the pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `awsElasticBlockStore` | `object` | Represents a Persistent Disk resource in AWS.<br><br>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling. |
| `azureDisk` | `object` | AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod. |
| `azureFile` | `object` | AzureFile represents an Azure File Service mount on the host and bind mount to the pod. |
| `cephfs` | `object` | Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling. |
| `cinder` | `object` | Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership |

| Property | Type | Description |
| --- | --- | --- |
| | | management and SELinux relabeling. |
| `configMap` | `object` | Adapts a ConfigMap into a volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling. |
| `csi` | `object` | Represents a source location of a volume to mount, managed by an external CSI driver |
| `downwardAPI` | `object` | DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling. |
| `emptyDir` | `object` | Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling. |
| `ephemeral` | `object` | Represents an ephemeral volume that is handled by a normal storage driver. |
| `fc` | `object` | Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling. |
| `flexVolume` | `object` | FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin. |
| `flocker` | `object` | Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling. |
| `gcePersistentDisk` | `object` | Represents a Persistent Disk resource in Google Compute Engine.<br><br>A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling. |
| `gitRepo` | `object` | Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.<br><br>DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container. |

| Property | Type | Description |
| --- | --- | --- |
| `glusterfs` | `object` | Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling. |
| `hostPath` | `object` | Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling. |
| `image` | `object` | ImageVolumeSource represents a image volume resource. |
| `iscsi` | `object` | Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling. |
| `name` | `string` | name of the volume. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `nfs` | `object` | Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling. |
| `persistentVolumeClaim` | `object` | PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system). |
| `photonPersistentDisk` | `object` | Represents a Photon Controller persistent disk resource. |
| `portworxVolume` | `object` | PortworxVolumeSource represents a Portworx volume resource. |
| `projected` | `object` | Represents a projected volume source |
| `quobyte` | `object` | Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling. |
| `rbd` | `object` | Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling. |
| `scaleIO` | `object` | ScaleIOVolumeSource represents a persistent ScaleIO volume |

| Property | Type | Description |
| --- | --- | --- |
| `secret` | `object` | Adapts a Secret into a volume.<br><br>The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling. |
| `storageos` | `object` | Represents a StorageOS persistent volume resource. |
| `vsphereVolume` | `object` | Represents a vSphere volume resource. |

## .spec.template.spec.volumes[].awsElasticBlockStore

**Description**

Represents a Persistent Disk resource in AWS. An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). |
| `readOnly` | `boolean` | readOnly value true will force the readOnly setting in VolumeMounts. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `volumeID` | `string` | volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |

## .spec.template.spec.volumes[].azureDisk

**Description**

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`diskName` `diskURI`

| Property | Type | Description |
|---|---|---|
| `cachingMode` | `string` | cachingMode is the Host Caching mode: None, Read Only, Read Write.<br>Possible enum values:<br>• `"None"`<br>• `"ReadOnly"`<br>• `"ReadWrite"` |
| `diskName` | `string` | diskName is the Name of the data disk in the blob storage |
| `diskURI` | `string` | diskURI is the URI of data disk in the blob storage |
| `fsType` | `string` | fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `kind` | `string` | kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared<br>Possible enum values:<br>• `"Dedicated"`<br>• `"Managed"`<br>• `"Shared"` |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

## .spec.template.spec.volumes[].azureFile

**Description**

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`secretName` `shareName`

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretName` | `string` | secretName is the name of secret that contains Azure Storage Account Name and Key |
| `shareName` | `string` | shareName is the azure share Name |

## .spec.template.spec.volumes[].cephfs

**Description**

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`monitors`

| Property | Type | Description |
| --- | --- | --- |
| `monitors` | `array` | monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `path` | `string` | path is Optional: Used as the mounted root, rather than the full Ceph tree, default is / |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretFile` | `string` | secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].cephfs.monitors

**Description**

monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].cephfs.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].cephfs.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].cinder

**Description**

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeID` | `string` | volumeID used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |

## .spec.template.spec.volumes[].cinder.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].configMap

**Description**

Adapts a ConfigMap into a volume. The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `defaultMode` | `integer` | defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].csi

**Description**

Represents a source location of a volume to mount, managed by an external CSI driver

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
|---|---|---|
| `driver` | `string` | driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster. |
| `fsType` | `string` | fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply. |

| Property | Type | Description |
| --- | --- | --- |
| `nodePublishSecretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `readOnly` | `boolean` | readOnly specifies a read-only configuration for the volume. Defaults to false (read/write). |
| `volumeAttributes` | `object` | volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values. |

## .spec.template.spec.volumes[].csi.nodePublishSecretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].csi.volumeAttributes

**Description**

volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**Type**

`object`

## .spec.template.spec.volumes[].downwardAPI

**Description**

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | Items is a list of downward API volume file |

## .spec.template.spec.volumes[].downwardAPI.items

**Description**

Items is a list of downward API volume file

**Type**

`array`

## .spec.template.spec.volumes[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
|---|---|---|
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].emptyDir

### Description

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `medium` | `string` | medium represents what type of storage medium should back this directory. The default is "" which means to use the n |
| `sizeLimit` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAM<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits> <r<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.html)<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a number<br>When a Quantity is parsed from a string, it will remember the type of suffix it had, an<br>Before serializing, Quantity will be put in "canonical form". This means that Exponent/<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suffix<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point number.<br>Non-canonical values will still parse as long as they are well formed, but will be re-e<br>This format is intended to make it difficult to use these numbers without writing some |

## .spec.template.spec.volumes[].ephemeral

**Description**

Represents an ephemeral volume that is handled by a normal storage driver.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `volumeClaimTemplate` | `object` | PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate

**Description**

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

**Type**

`object`

**Required**

`spec`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec

### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `accessModes` | `array` | accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1 ↗ |
| `dataSource` | `object` | TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace. |
| `dataSourceRef` | `object` | TypedObjectReference contains enough information to let you locate the typed referenced object |
| `resources` | `object` | VolumeResourceRequirements describes the storage resource requirements for a volume. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `storageClassName` | `string` | storageClassName is the name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1 ↗ |
| `volumeAttributesClassName` | `string` | volumeAttributesClassName may be used to set the VolumeAttributesClass used by this claim. If specified, the CSI driver will create or update the volume with the attributes defined in the corresponding VolumeAttributesClass. This has a different purpose than storageClassName, it can be changed after the claim is created. An empty string value means that no VolumeAttributesClass will be applied to the claim but it's not allowed to reset this field to empty string once it is set. If unspecified and the PersistentVolumeClaim is unbound, the default VolumeAttributesClass will be set by the persistentvolume controller if it exists. If the resource referred to by volumeAttributesClass does not exist, this PersistentVolumeClaim will be set to a Pending state, as reflected by the |

| Property | Type | Description |
| --- | --- | --- |
| | | modifyVolumeStatus field, until such as a resource exists. More info: https://kubernetes.io/docs/concepts/storage/volume-attributes-classes/ (Beta) Using this field requires the VolumeAttributesClass feature gate to be enabled (off by default). |
| `volumeMode` | `string` | volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.<br><br>Possible enum values:<br>• `"Block"` means the volume will not be formatted with a filesystem and will remain a raw block device.<br>• `"Filesystem"` means the volume will be or is formatted with a filesystem. |
| `volumeName` | `string` | volumeName is the binding reference to the PersistentVolume backing this claim. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes

**Description**

accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

**Description**

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is the name of resource being referenced |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef

**Description**

TypedObjectReference contains enough information to let you locate the typed referenced object

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |
| `namespace` | `string` | Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

**Description**

VolumeResourceRequirements describes the storage resource requirements for a volume.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].fc

**Description**

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `lun` | `integer` | lun is Optional: FC target lun number |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `targetWWNs` | `array` | targetWWNs is Optional: FC target worldwide names (WWNs) |
| `wwids` | `array` | wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously. |

## .spec.template.spec.volumes[].fc.targetWWNs

**Description**

targetWWNs is Optional: FC target worldwide names (WWNs)

**Type**

`array`

## .spec.template.spec.volumes[].fc.targetWWNs[]

**Type**

`string`

## .spec.template.spec.volumes[].fc.wwids

**Description**

wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

**Type**

`array`

## .spec.template.spec.volumes[].fc.wwids[]

**Type**

`string`

## .spec.template.spec.volumes[].flexVolume

**Description**

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
| --- | --- | --- |
| `driver` | `string` | driver is the name of the driver to use for this volume. |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script. |
| `options` | `object` | options is Optional: this field holds extra command options if any. |
| `readOnly` | `boolean` | readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |

## .spec.template.spec.volumes[].flexVolume.options

**Description**

options is Optional: this field holds extra command options if any.

**Type**

`object`

## .spec.template.spec.volumes[].flexVolume.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: |

| Property | Type | Description |
| --- | --- | --- |
| | | https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].flocker

### Description

Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `datasetName` | `string` | datasetName is Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated |
| `datasetUUID` | `string` | datasetUUID is the UUID of the dataset. This is unique identifier of a Flocker dataset |

## .spec.template.spec.volumes[].gcePersistentDisk

### Description

Represents a Persistent Disk resource in Google Compute Engine. A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

### Type

`object`

### Required

`pdName`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `pdName` | `string` | pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

## .spec.template.spec.volumes[].gitRepo

**Description**

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

**Type**

`object`

**Required**

`repository`

| Property | Type | Description |
| --- | --- | --- |
| `directory` | `string` | directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name. |
| `repository` | `string` | repository is the URL |
| `revision` | `string` | revision is the commit hash for the specified revision. |

## .spec.template.spec.volumes[].glusterfs

**Description**

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`endpoints` `path`

| Property | Type | Description |
| --- | --- | --- |
| `endpoints` | `string` | endpoints is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |
| `path` | `string` | path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

## .spec.template.spec.volumes[].hostPath

### Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`path`

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗ |
| `type` | `string` | type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗<br><br>Possible enum values:<br>• `""` For backwards compatible, leave it empty if unset<br>• `"BlockDevice"` A block device must exist at the given path<br>• `"CharDevice"` A character device must exist at the given path<br>• `"Directory"` A directory must exist at the given path<br>• `"DirectoryOrCreate"` If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet.<br>• `"File"` A file must exist at the given path<br>• `"FileOrCreate"` If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet.<br>• `"Socket"` A UNIX socket must exist at the given path |

## .spec.template.spec.volumes[].image

### Description

ImageVolumeSource represents a image volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `pullPolicy` | `string` | Policy for pulling OCI objects. Possible values are: Always: the kubelet always attempts to pull the reference. Container creation will fail If the pull fails. Never: the kubelet never pulls the reference and only uses a local image or artifact. |

| Property | Type | Description |
| --- | --- | --- |
| | | Container creation will fail if the reference isn't present. IfNotPresent: the kubelet pulls if the reference isn't already present on disk. Container creation will fail if the reference isn't present and the pull fails. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise.<br><br>Possible enum values:<br><br>- `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>- `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>- `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `reference` | `string` | Required: Image or artifact reference to be used. Behaves in the same way as pod.spec.containers[*].image. Pull secrets will be assembled in the same way as for the container image by looking up node credentials, SA image pull secrets, and pod spec image pull secrets. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |

## .spec.template.spec.volumes[].iscsi

### Description

Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`targetPortal` `iqn` `lun`

| Property | Type | Description |
| --- | --- | --- |
| `chapAuthDiscovery` | `boolean` | chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication |
| `chapAuthSession` | `boolean` | chapAuthSession defines whether support iSCSI Session CHAP authentication |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi |
| `initiatorName` | `string` | initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection. |
| `iqn` | `string` | iqn is the target iSCSI Qualified Name. |

| Property | Type | Description |
| --- | --- | --- |
| `iscsiInterface` | `string` | iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp). |
| `lun` | `integer` | lun represents iSCSI Target Lun number. |
| `portals` | `array` | portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `targetPortal` | `string` | targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |

## .spec.template.spec.volumes[].iscsi.portals

**Description**

portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

**Type**

`array`

## .spec.template.spec.volumes[].iscsi.portals[]

**Type**

`string`

## .spec.template.spec.volumes[].iscsi.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].nfs

**Description**

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`server` `path`

| Property | Type | Description |
|---|---|---|
| `path` | `string` | path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |
| `readOnly` | `boolean` | readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |
| `server` | `string` | server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |

## .spec.template.spec.volumes[].persistentVolumeClaim

**Description**

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

**Type**

`object`

**Required**

`claimName`

| Property | Type | Description |
|---|---|---|
| `claimName` | `string` | claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims |
| `readOnly` | `boolean` | readOnly Will force the ReadOnly setting in VolumeMounts. Default false. |

## .spec.template.spec.volumes[].photonPersistentDisk

**Description**

Represents a Photon Controller persistent disk resource.

**Type**

`object`

**Required**

`pdID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `pdID` | `string` | pdID is the ID that identifies Photon Controller persistent disk |

## .spec.template.spec.volumes[].portworxVolume

**Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `volumeID` | `string` | volumeID uniquely identifies a Portworx volume |

## .spec.template.spec.volumes[].projected

**Description**

Represents a projected volume source

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `sources` | `array` | sources is the list of volume projections. Each entry in this list handles one source. |

## .spec.template.spec.volumes[].projected.sources

**Description**

sources is the list of volume projections. Each entry in this list handles one source.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[]

**Description**

Projection that may be projected along with other supported volume types. Exactly one of these fields must be set.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `clusterTrustBundle` | `object` | ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem. |
| `configMap` | `object` | Adapts a ConfigMap into a projected volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode. |
| `downwardAPI` | `object` | Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode. |
| `secret` | `object` | Adapts a secret into a projected volume.<br><br>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode. |
| `serviceAccountToken` | `object` | ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise). |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle

**Description**

ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem.

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `name` | `string` | Select a single ClusterTrustBundle by object name. Mutually-exclusive with signerName and labelSelector. |
| `optional` | `boolean` | If true, don't block pod startup if the referenced ClusterTrustBundle(s) aren't available. If using name, then the named ClusterTrustBundle is allowed not to exist. If using signerName, then the combination of signerName and labelSelector is allowed to match zero ClusterTrustBundles. |
| `path` | `string` | Relative path from the volume root to write the bundle. |
| `signerName` | `string` | Select all ClusterTrustBundles that match this signer name. Mutually-exclusive with name. The contents of all selected ClusterTrustBundles will be unified and deduplicated. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].projected.sources[].configMap

**Description**

Adapts a ConfigMap into a projected volume. The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].projected.sources[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI

**Description**

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | Items is a list of DownwardAPIVolume file |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items

**Description**

Items is a list of DownwardAPIVolume file

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |

| Property | Type | Description |
| --- | --- | --- |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and |

| Property | Type | Description |
| --- | --- | --- |
| | | The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].projected.sources[].secret

### Description

Adapts a secret into a projected volume. The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | optional field specify whether the Secret or its key must be defined |

## .spec.template.spec.volumes[].projected.sources[].secret.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].serviceAccountToken

**Description**

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `audience` | `string` | audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver. |
| `expirationSeconds` | `integer` | expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours.Defaults to 1 hour and must be at least 10 minutes. |
| `path` | `string` | path is the path relative to the mount point of the file to project the token into. |

## .spec.template.spec.volumes[].quobyte

**Description**

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`registry` `volume`

| Property | Type | Description |
|---|---|---|
| `group` | `string` | group to map volume access to Default is no group |
| `readOnly` | `boolean` | readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false. |
| `registry` | `string` | registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes |
| `tenant` | `string` | tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin |
| `user` | `string` | user to map volume access to Defaults to serivceaccount user |

| Property | Type | Description |
| --- | --- | --- |
| `volume` | `string` | volume is a string that references an already created Quobyte volume by name. |

## .spec.template.spec.volumes[].rbd

### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`monitors` `image`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd ↗ |
| `image` | `string` | image is the rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `keyring` | `string` | keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `monitors` | `array` | monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `pool` | `string` | pool is the rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is the rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].rbd.monitors

**Description**

monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].rbd.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].rbd.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].scaleIO

**Description**

ScaleIOVolumeSource represents a persistent ScaleIO volume

**Type**

`object`

**Required**

`gateway` `system` `secretRef`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs". |
| `gateway` | `string` | gateway is the host address of the ScaleIO API Gateway. |
| `protectionDomain` | `string` | protectionDomain is the name of the ScaleIO Protection Domain for the configured storage. |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

| Property | Type | Description |
| --- | --- | --- |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `sslEnabled` | `boolean` | sslEnabled Flag enable/disable SSL communication with Gateway, default false |
| `storageMode` | `string` | storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned. |
| `storagePool` | `string` | storagePool is the ScaleIO Storage Pool associated with the protection domain. |
| `system` | `string` | system is the name of the storage system as configured in ScaleIO. |
| `volumeName` | `string` | volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source. |

## .spec.template.spec.volumes[].scaleIO.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].secret

**Description**

Adapts a Secret into a volume. The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON |

| Property | Type | Description |
| --- | --- | --- |
| | | requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `optional` | `boolean` | optional field specify whether the Secret or its keys must be defined |
| `secretName` | `string` | secretName is the name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret ↗ |

## .spec.template.spec.volumes[].secret.items

**Description**

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].storageos

### Description

Represents a StorageOS persistent volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeName` | `string` | volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace. |
| `volumeNamespace` | `string` | volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created. |

## .spec.template.spec.volumes[].storageos.secretRef

### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].vsphereVolume

**Description**

Represents a vSphere volume resource.

**Type**

`object`

**Required**

`volumePath`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `storagePolicyID` | `string` | storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName. |
| `storagePolicyName` | `string` | storagePolicyName is the storage Policy Based Management (SPBM) profile name. |
| `volumePath` | `string` | volumePath is the path that identifies vSphere volume vmdk |

## .status

**Description**

DeploymentStatus is the most recently observed status of the Deployment.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `availableReplicas` | `integer` | Total number of available pods (ready for at least minReadySeconds) targeted by this deployment. |
| `collisionCount` | `integer` | Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet. |
| `conditions` | `array` | Represents the latest available observations of a deployment's current state. |
| `observedGeneration` | `integer` | The generation observed by the deployment controller. |
| `readyReplicas` | `integer` | readyReplicas is the number of pods targeted by this Deployment with a Ready Condition. |

| Property | Type | Description |
| --- | --- | --- |
| `replicas` | `integer` | Total number of non-terminated pods targeted by this deployment (their labels match the selector). |
| `unavailableReplicas` | `integer` | Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created. |
| `updatedReplicas` | `integer` | Total number of non-terminated pods targeted by this deployment that have the desired template spec. |

## .status.conditions

**Description**

Represents the latest available observations of a deployment's current state.

**Type**

`array`

## .status.conditions[]

**Description**

DeploymentCondition describes the state of a deployment at a certain point.

**Type**

`object`

**Required**

`type` `status`

| Property | Type | Description |
| --- | --- | --- |
| `lastTransitionTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `lastUpdateTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | A human readable message indicating details about the transition. |
| `reason` | `string` | The reason for the condition's last transition. |
| `status` | `string` | Status of the condition, one of True, False, Unknown. |

| Property | Type | Description |
| --- | --- | --- |
| `type` | `string` | Type of deployment condition. |

# API Endpoints

The following API endpoints are available:

- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/deployments`
  - `DELETE` : delete collection of Deployment
  - `GET` : list objects of kind Deployment
  - `POST` : create a new Deployment
- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/deployments/{name}`
  - `DELETE` : delete the specified Deployment
  - `GET` : read the specified Deployment
  - `PATCH` : partially update the specified Deployment
  - `PUT` : replace the specified Deployment
- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/deployments/{name}/status`
  - `GET` : read status of the specified Deployment
  - `PATCH` : partially update status of the specified Deployment
  - `PUT` : replace status of the specified Deployment

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/deployments

**HTTP method**

`DELETE`

**Description**

delete collection of Deployment

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`GET`

**Description**

list objects of kind Deployment

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `DeploymentList` schema |

| HTTP code | Response body |
| --- | --- |
| 401 - Unauthorized | Empty |

### HTTP method

`POST`

### Description

create a new Deployment

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `Deployment` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Deployment` schema |
| 201 - Created | `Deployment` schema |
| 202 - Accepted | `Deployment` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/deployments/{name}

### HTTP method

`DELETE`

### Description

delete the specified Deployment

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 202 - Accepted | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

### Description

read the specified Deployment

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Deployment` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update the specified Deployment

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Deployment` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace the specified Deployment

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
|---|---|---|
| `body` | `Deployment` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `Deployment` schema |
| 201 - Created | `Deployment` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/deployments/{name}/status

### HTTP method

`GET`

### Description

read status of the specified Deployment

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `Deployment` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update status of the specified Deployment

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `Deployment` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace status of the specified Deployment

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
|---|---|---|
| `body` | `Deployment` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Deployment` schema |
| 201 - Created | `Deployment` schema |
| 401 - Unauthorized | Empty |

# Job [batch/v1]

**Description**

Job represents the configuration of a single job.

**Type**

`object`

## Specification

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources |
| `kind` | `string` | Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | JobSpec describes how the job execution will look like. |
| `status` | `object` | JobStatus represents the current state of a Job. |

### .spec

**Description**

JobSpec describes how the job execution will look like.

**Type**

`object`

**Required**

`template`

| Property | Type | Description |
| --- | --- | --- |
| `activeDeadlineSeconds` | `integer` | Specifies the duration in seconds relative to the startTime that the job may be continuously active before the system tries to terminate it; value must be positive integer. If a Job is suspended (at creation or through an update), this timer will effectively be stopped and reset when the Job is resumed again. |
| `backoffLimit` | `integer` | Specifies the number of retries before marking this job failed. Defaults to 6 |
| `backoffLimitPerIndex` | `integer` | Specifies the limit for the number of retries within an index before marking this index as failed. When enabled the number of failures per index is kept in the pod's batch.kubernetes.io/job-index-failure-count annotation. It can only be set when Job's completionMode=Indexed, and the Pod's restart policy is Never. The field is immutable. This field is beta-level. It can be used when the `JobBackoffLimitPerIndex` feature gate is enabled (enabled by default). |
| `completionMode` | `string` | completionMode specifies how Pod completions are tracked. It can be `NonIndexed` (default) or `Indexed`.<br><br>`NonIndexed` means that the Job is considered complete when there have been .spec.completions successfully completed Pods. Each Pod completion is homologous to each other.<br><br>`Indexed` means that the Pods of a Job get an associated completion index from 0 to (.spec.completions - 1), available in the annotation batch.kubernetes.io/job-completion-index. The Job is considered complete when there is one successfully completed Pod for each index. When value is `Indexed`, .spec.completions must be specified and `.spec.parallelism` must be less than or equal to 10^5. In addition, The Pod name takes the form `$(job-name)-$(index)-$(random-string)`, the Pod hostname takes the form `$(job-name)-$(index)`.<br><br>More completion modes can be added in the future. If the Job controller observes a mode that it doesn't recognize, which is possible during upgrades due to version skew, the controller skips updates for the Job.<br><br>Possible enum values:<br><br>- `"Indexed"` is a Job completion mode. In this mode, the Pods of a Job get an associated completion index from 0 to (.spec.completions - 1). The Job is considered complete when a Pod completes for each completion index.<br>- `"NonIndexed"` is a Job completion mode. In this mode, the Job is considered complete when there have been .spec.completions successfully completed Pods. Pod completions are homologous to each other. |
| `completions` | `integer` | Specifies the desired number of successfully finished pods the job should be run with. Setting to null means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/ |
| `managedBy` | `string` | ManagedBy field indicates the controller that manages a Job. The k8s Job controller reconciles jobs which don't have this field at all or the field value is the reserved string `kubernetes.io/job-` |

| Property | Type | Description |
| --- | --- | --- |
| | | `controller`, but skips reconciling Jobs with a custom value for this field. The value must be a valid domain-prefixed path (e.g. acme.io/foo) - all characters before the first "/" must be a valid subdomain as defined by RFC 1123. All characters trailing the first "/" must be valid HTTP Path characters as defined by RFC 3986. The value cannot exceed 63 characters. This field is immutable.<br><br>This field is beta-level. The job controller accepts setting the field when the feature gate JobManagedBy is enabled (enabled by default). |
| `manualSelector` | `boolean` | manualSelector controls generation of pod labels and pod selectors. Leave `manualSelector` unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see `manualSelector=true` in jobs that were created with the old `extensions/v1beta1` API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector ↗ |
| `maxFailedIndexes` | `integer` | Specifies the maximal number of failed indexes before marking the Job as failed, when backoffLimitPerIndex is set. Once the number of failed indexes exceeds this number the entire Job is marked as Failed and its execution is terminated. When left as null the job continues execution of all of its indexes and is marked with the `Complete` Job condition. It can only be specified when backoffLimitPerIndex is set. It can be null or up to completions. It is required and must be less than or equal to 10^4 when is completions greater than 10^5. This field is beta-level. It can be used when the `JobBackoffLimitPerIndex` feature gate is enabled (enabled by default). |
| `parallelism` | `integer` | Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when ((.spec.completions - .status.successful) < .spec.parallelism), i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/ ↗ |
| `podFailurePolicy` | `object` | PodFailurePolicy describes how failed pods influence the backoffLimit. |
| `podReplacementPolicy` | `string` | podReplacementPolicy specifies when to create replacement Pods. Possible values are: - TerminatingOrFailed means that we recreate pods when they are terminating (has a metadata.deletionTimestamp) or failed.<br><br>• Failed means to wait until a previously created Pod is fully terminated (has phase Failed or Succeeded) before creating a replacement Pod.<br><br>When using podFailurePolicy, Failed is the the only allowed value. TerminatingOrFailed and Failed are allowed values when podFailurePolicy is not in use. This is an beta field. To use this, enable the JobPodReplacementPolicy feature toggle. This is on by default.<br><br>Possible enum values:<br><br>• `"Failed"` means to wait until a previously created Pod is fully terminated (has phase Failed or Succeeded) before creating a replacement Pod. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"TerminatingOrFailed"` means that we recreate pods when they are terminating (has a metadata.deletionTimestamp) or failed. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `successPolicy` | `object` | SuccessPolicy describes when a Job can be declared as succeeded based on the success of some indexes. |
| `suspend` | `boolean` | suspend specifies whether the Job controller should create Pods or not. If a Job is created with suspend set to true, no Pods are created by the Job controller. If a Job is suspended after creation (i.e. the flag goes from false to true), the Job controller will delete all active Pods associated with this Job. Users must design their workload to gracefully handle this. Suspending a Job will reset the StartTime field of the Job, effectively resetting the ActiveDeadlineSeconds timer too. Defaults to false. |
| `template` | `object` | PodTemplateSpec describes the data a pod should have when created from a template |
| `ttlSecondsAfterFinished` | `integer` | ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. |

## .spec.podFailurePolicy

**Description**

PodFailurePolicy describes how failed pods influence the backoffLimit.

**Type**

`object`

**Required**

`rules`

| Property | Type | Description |
| --- | --- | --- |
| `rules` | `array` | A list of pod failure policy rules. The rules are evaluated in order. Once a rule matches a Pod failure, the remaining of the rules are ignored. When no rule matches the Pod failure, the default handling applies - the counter of pod failures is incremented and it is checked against the backoffLimit. At most 20 elements are allowed. |

## .spec.podFailurePolicy.rules

**Description**

A list of pod failure policy rules. The rules are evaluated in order. Once a rule matches a Pod failure, the remaining of the rules are ignored. When no rule matches the Pod failure, the default handling applies - the counter of pod failures is incremented and it is checked against the backoffLimit. At most 20 elements are allowed.

**Type**

`array`

## .spec.podFailurePolicy.rules[]

**Description**

PodFailurePolicyRule describes how a pod failure is handled when the requirements are met. One of onExitCodes and onPodConditions, but not both, can be used in each rule.

**Type**

`object`

**Required**

`action`

| Property | Type | Description |
| --- | --- | --- |
| `action` | `string` | Specifies the action taken on a pod failure when the requirements are satisfied. Possible values are:<br>• FailJob: indicates that the pod's job is marked as Failed and all running pods are terminated.<br>• FailIndex: indicates that the pod's index is marked as Failed and will not be restarted. This value is beta-level. It can be used when the `JobBackoffLimitPerIndex` feature gate is enabled (enabled by default).<br>• Ignore: indicates that the counter towards the .backoffLimit is not incremented and a replacement pod is created.<br>• Count: indicates that the pod is handled in the default way - the counter towards the .backoffLimit is incremented. Additional values are considered to be added in the future. Clients should react to an unknown action by skipping the rule.<br>Possible enum values:<br>• `"Count"` This is an action which might be taken on a pod failure - the pod failure is handled in the default way - the counter towards .backoffLimit, represented by the job's .status.failed field, is incremented.<br>• `"FailIndex"` This is an action which might be taken on a pod failure - mark the Job's index as failed to avoid restarts within this index. This action can only be used when backoffLimitPerIndex is set. This value is beta-level.<br>• `"FailJob"` This is an action which might be taken on a pod failure - mark the pod's job as Failed and terminate all running pods.<br>• `"Ignore"` This is an action which might be taken on a pod failure - the counter towards .backoffLimit, represented by the job's .status.failed field, is not incremented and a replacement pod is created. |
| `onExitCodes` | `object` | PodFailurePolicyOnExitCodesRequirement describes the requirement for handling a failed pod based on its container exit codes. In particular, it lookups the .state.terminated.exitCode for each app container and init container status, represented by the .status.containerStatuses and .status.initContainerStatuses fields in the Pod status, respectively. Containers completed with success (exit code 0) are excluded from the requirement check. |

| Property | Type | Description |
| --- | --- | --- |
| `onPodConditions` | `array` | Represents the requirement on the pod conditions. The requirement is represented as a list of pod condition patterns. The requirement is satisfied if at least one pattern matches an actual pod condition. At most 20 elements are allowed. |

## .spec.podFailurePolicy.rules[].onExitCodes

**Description**

PodFailurePolicyOnExitCodesRequirement describes the requirement for handling a failed pod based on its container exit codes. In particular, it lookups the .state.terminated.exitCode for each app container and init container status, represented by the .status.containerStatuses and .status.initContainerStatuses fields in the Pod status, respectively. Containers completed with success (exit code 0) are excluded from the requirement check.

**Type**

`object`

**Required**

`operator` `values`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Restricts the check for exit codes to the container with the specified name. When null, the rule applies to all containers. When specified, it should match one the container or initContainer names in the pod template. |
| `operator` | `string` | Represents the relationship between the container exit code(s) and the specified values. Containers completed with success (exit code 0) are excluded from the requirement check. Possible values are:<br>• In: the requirement is satisfied if at least one container exit code (might be multiple if there are multiple containers not restricted by the 'containerName' field) is in the set of specified values.<br>• NotIn: the requirement is satisfied if at least one container exit code (might be multiple if there are multiple containers not restricted by the 'containerName' field) is not in the set of specified values. Additional values are considered to be added in the future. Clients should react to an unknown operator by assuming the requirement is not satisfied.<br>Possible enum values:<br>• `"In"`<br>• `"NotIn"` |
| `values` | `array` | Specifies the set of values. Each returned container exit code (might be multiple in case of multiple containers) is checked against this set of values with respect to the operator. The list of values must be ordered and must not contain duplicates. Value '0' cannot be used for the In operator. At least one element is required. At most 255 elements are allowed. |

## .spec.podFailurePolicy.rules[].onExitCodes.values

**Description**

Specifies the set of values. Each returned container exit code (might be multiple in case of multiple containers) is checked against this set of values with respect to the operator. The list of values must be ordered and must not contain duplicates. Value '0' cannot be used for the In operator. At least one element is required. At most 255 elements are allowed.

**Type**

`array`

## .spec.podFailurePolicy.rules[].onExitCodes.values[]

**Type**

`integer`

## .spec.podFailurePolicy.rules[].onPodConditions

**Description**

Represents the requirement on the pod conditions. The requirement is represented as a list of pod condition patterns. The requirement is satisfied if at least one pattern matches an actual pod condition. At most 20 elements are allowed.

**Type**

`array`

## .spec.podFailurePolicy.rules[].onPodConditions[]

**Description**

PodFailurePolicyOnPodConditionsPattern describes a pattern for matching an actual pod condition type.

**Type**

`object`

**Required**

`type` `status`

| Property | Type | Description |
| --- | --- | --- |
| `status` | `string` | Specifies the required Pod condition status. To match a pod condition it is required that the specified status equals the pod condition status. Defaults to True. |
| `type` | `string` | Specifies the required Pod condition type. To match a pod condition it is required that specified type equals the pod condition type. |

## .spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.successPolicy

**Description**

SuccessPolicy describes when a Job can be declared as succeeded based on the success of some indexes.

**Type**

`object`

**Required**

`rules`

| Property | Type | Description |
| --- | --- | --- |
| `rules` | `array` | rules represents the list of alternative rules for the declaring the Jobs as successful before `.status.succeeded >= .spec.completions`. Once any of the rules are met, the "SucceededCriteriaMet" condition is added, and the lingering pods are removed. The terminal state for such a Job has the "Complete" condition. Additionally, these rules are evaluated in order; Once the Job meets one of the rules, other rules are ignored. At most 20 elements are allowed. |

## .spec.successPolicy.rules

**Description**

rules represents the list of alternative rules for the declaring the Jobs as successful before `.status.succeeded >= .spec.completions`. Once any of the rules are met, the "SucceededCriteriaMet" condition is added, and the lingering pods are removed. The terminal state for such a Job has the "Complete" condition. Additionally, these rules are evaluated in order; Once the Job meets one of the rules, other rules are ignored. At most 20 elements are allowed.

**Type**

`array`

## .spec.successPolicy.rules[]

**Description**

SuccessPolicyRule describes rule for declaring a Job as succeeded. Each rule must have at least one of the "succeededIndexes" or "succeededCount" specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `succeededCount` | `integer` | succeededCount specifies the minimal required size of the actual set of the succeeded indexes for the Job. When succeededCount is used along with succeededIndexes, the check is constrained only to the set of indexes specified by succeededIndexes. For example, given that succeededIndexes is "1-4", succeededCount is "3", and completed indexes are "1", "3", and "5", the Job isn't declared as succeeded because only "1" and "3" indexes are considered in that rules. When this field is null, this doesn't default to any value and is never evaluated at any time. When specified it needs to be a positive integer. |
| `succeededIndexes` | `string` | succeededIndexes specifies the set of indexes which need to be contained in the actual set of the succeeded indexes for the Job. The list of indexes must be within 0 to ".spec.completions-1" and must not contain duplicates. At least one element is required. The indexes are represented as intervals separated by commas. The intervals can be a decimal integer or a pair of decimal integers separated by a hyphen. The number are listed in represented by the first and last element of the series, separated by a hyphen. For example, if the completed indexes are 1, 3, 4, 5 and 7, they are represented as "1,3-5,7". When this field is null, this field doesn't default to any value and is never evaluated at any time. |

## .spec.template

### Description

PodTemplateSpec describes the data a pod should have when created from a template

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PodSpec is a description of a pod. |

## .spec.template.spec

### Description

PodSpec is a description of a pod.

### Type

`object`

### Required

`containers`

| Property | Type | Description |
| --- | --- | --- |
| `activeDeadlineSeconds` | `integer` | Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer. |

| Property | Type | Description |
| --- | --- | --- |
| `affinity` | `object` | Affinity is a group of affinity scheduling rules. |
| `automountServiceAccountToken` | `boolean` | AutomountServiceAccountToken indicates whether a service account token should be automatically mounted. |
| `containers` | `array` | List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated. |
| `dnsConfig` | `object` | PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy. |
| `dnsPolicy` | `string` | Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.<br><br>Possible enum values:<br>• `"ClusterFirst"` indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>• `"ClusterFirstWithHostNet"` indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>• `"Default"` indicates that the pod should use the default (as determined by kubelet) DNS settings.<br>• `"None"` indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig. |
| `enableServiceLinks` | `boolean` | EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching traditional container linking syntax. Optional: Defaults to true. |
| `ephemeralContainers` | `array` | List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. |
| `hostAliases` | `array` | HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. |
| `hostIPC` | `boolean` | Use the host's ipc namespace. Optional: Default to false. |

| Property | Type | Description |
| --- | --- | --- |
| `hostNetwork` | `boolean` | Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false. |
| `hostPID` | `boolean` | Use the host's pid namespace. Optional: Default to false. |
| `hostUsers` | `boolean` | Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new userns is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature. |
| `hostname` | `string` | Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value. |
| `imagePullSecrets` | `array` | ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod ↗ |
| `initContainers` | `array` | List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/ ↗ |
| `nodeName` | `string` | NodeName indicates in which node this pod is scheduled. If empty, this pod is a candidate for scheduling by the scheduler defined in schedulerName. Once this field is set, the kubelet for this node becomes responsible for the lifecycle of this pod. This field should not be used to express a desire for the pod to be scheduled on a specific node. https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodename ↗ |
| `nodeSelector` | `object` | NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/ ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `os` | `object` | PodOS defines the OS parameters of a pod. |
| `overhead` | `object` | Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md ↗ |
| `preemptionPolicy` | `string` | PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.<br><br>Possible enum values:<br>• `"Never"` means that pod never preempts other pods with lower priority.<br>• `"PreemptLowerPriority"` means that pod can preempt other pods with lower priority. |
| `priority` | `integer` | The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority. |
| `priorityClassName` | `string` | If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default. |
| `readinessGates` | `array` | If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates ↗ |
| `resourceClaims` | `array` | ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.<br><br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |

| Property | Type | Description |
| --- | --- | --- |
| `restartPolicy` | `string` | Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy ↗<br><br>Possible enum values:<br>• `"Always"`<br>• `"Never"`<br>• `"OnFailure"` |
| `runtimeClassName` | `string` | RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class ↗ |
| `schedulerName` | `string` | If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler. |
| `schedulingGates` | `array` | SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.<br><br>SchedulingGates can only be set at pod creation time, and be removed only afterwards. |
| `securityContext` | `object` | PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext. |
| `serviceAccount` | `string` | DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead. |
| `serviceAccountName` | `string` | ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/ ↗ |
| `setHostnameAsFQDN` | `boolean` | If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false. |

| Property | Type | Description |
| --- | --- | --- |
| `shareProcessNamespace` | `boolean` | Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false. |
| `subdomain` | `string` | If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all. |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds. |
| `tolerations` | `array` | If specified, the pod's tolerations. |
| `topologySpreadConstraints` | `array` | TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed. |
| `volumes` | `array` | List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes ↗ |

## .spec.template.spec.affinity

### Description

Affinity is a group of affinity scheduling rules.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `nodeAffinity` | `object` | Node affinity is a group of node affinity scheduling rules. |
| `podAffinity` | `object` | Pod affinity is a group of inter pod affinity scheduling rules. |
| `podAntiAffinity` | `object` | Pod anti affinity is a group of inter pod anti affinity scheduling rules. |

## .spec.template.spec.affinity.nodeAffinity

**Description**

Node affinity is a group of node affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `object` | A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

**Type**

`object`

**Required**

`weight` `preference`

| Property | Type | Description |
| --- | --- | --- |
| `preference` | `object` | A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm. |
| `weight` | `integer` | Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

**Type**

`object`

**Required**

`nodeSelectorTerms`

| Property | Type | Description |
|---|---|---|
| `nodeSelectorTerms` | `array` | Required. A list of node selector terms. The terms are ORed. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms

**Description**

Required. A list of node selector terms. The terms are ORed.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer.

This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity

**Description**

Pod affinity is a group of inter pod affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
|---|---|---|
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with |

| Property | Type | Description |
| --- | --- | --- |
| | | `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity

**Description**

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
|---|---|---|
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with |

| Property | Type | Description |
| --- | --- | --- |
| | | `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.containers

**Description**

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[]

**Description**

A single application container that you want to run within a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |

| Property | Type | Description |
| --- | --- | --- |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |

| Property | Type | Description |
| --- | --- | --- |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values: |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.containers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].args[]

**Type**

`string`

## .spec.template.spec.containers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].command[]

**Type**

`string`

## .spec.template.spec.containers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.containers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.containers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.containers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |

| Property | Type | Description |
| --- | --- | --- |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.containers[].env[].valueFrom.resourceFieldRef

### Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

### Type

`object`

### Required

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |

| Property | Type | Description |
| --- | --- | --- |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.containers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.containers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.containers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty.<br>Instances of this type with an empty value here are almost certainly wrong. More info:<br>https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.containers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty.<br>Instances of this type with an empty value here are almost certainly wrong. More info:<br>https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.containers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.containers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br><br>• `"HTTP"` means that the scheme used will be http:// |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
|---|---|---|
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.containers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.containers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

# .spec.template.spec.containers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

# .spec.template.spec.containers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

# .spec.template.spec.containers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |

| Property | Type | Description |
| --- | --- | --- |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.containers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>- `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>- `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.containers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>- `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>- `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.containers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.containers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.containers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are: |

| Property | Type | Description |
| --- | --- | --- |
| | | Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>- `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>- `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.containers[].securityContext.windowsOptions

### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.containers[].startupProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].startupProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].startupProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.containers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

# .spec.template.spec.containers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

# .spec.template.spec.containers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
|---|---|---|
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field |

| Property | Type | Description |
| --- | --- | --- |
| | | is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.dnsConfig

**Description**

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `nameservers` | `array` | A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed. |
| `options` | `array` | A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy. |
| `searches` | `array` | A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed. |

## .spec.template.spec.dnsConfig.nameservers

**Description**

A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.nameservers[]

**Type**

`string`

## .spec.template.spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

`array`

## .spec.template.spec.dnsConfig.options[]

**Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is this DNS resolver option's name. Required. |
| `value` | `string` | Value is this DNS resolver option's value. |

## .spec.template.spec.dnsConfig.searches

**Description**

A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.searches[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers

**Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[]

**Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `args` | `array` | Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers. |
| `ports` | `array` | Ports are not allowed for ephemeral containers. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for the container to manage the restart behavior of each container within a pod. This may only be set for init containers. You cannot set this field on ephemeral containers. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to |

| Property | Type | Description |
| --- | --- | --- |
| | | stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `targetContainerName` | `string` | If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.<br><br>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined. |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values:<br><br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.ephemeralContainers[].args

**Description**

Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].args[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.ephemeralContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.ephemeralContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].ports

**Description**

Ports are not allowed for ephemeral containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |

| Property | Type | Description |
| --- | --- | --- |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br><br>Possible enum values:<br>- `"SCTP"` is the SCTP protocol.<br>- `"TCP"` is the TCP protocol.<br>- `"UDP"` is the UDP protocol. |

## .spec.template.spec.ephemeralContainers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.ephemeralContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.ephemeralContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.ephemeralContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.ephemeralContainers[].securityContext.appArmorProfile

### Description

AppArmorProfile defines a pod or container's AppArmor settings.

### Type

`object`

### Required

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.ephemeralContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.ephemeralContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.ephemeralContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.ephemeralContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.hostAliases

**Description**

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified.

**Type**

`array`

## .spec.template.spec.hostAliases[]

**Description**

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

**Type**

`object`

**Required**

`ip`

| Property | Type | Description |
|---|---|---|
| `hostnames` | `array` | Hostnames for the above IP address. |
| `ip` | `string` | IP address of the host file entry. |

## .spec.template.spec.hostAliases[].hostnames

**Description**

Hostnames for the above IP address.

**Type**

`array`

## .spec.template.spec.hostAliases[].hostnames[]

**Type**

`string`

## .spec.template.spec.imagePullSecrets

**Description**

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

**Type**

`array`

## .spec.template.spec.imagePullSecrets[]

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.initContainers

### Description

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

### Type

`array`

## .spec.template.spec.initContainers[]

### Description

A single application container that you want to run within a pod.

### Type

`object`

### Required

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |

| Property | Type | Description |
| --- | --- | --- |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.initContainers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].args[]

**Type**

`string`

## .spec.template.spec.initContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].command[]

**Type**

`string`

## .spec.template.spec.initContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

# .spec.template.spec.initContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

# .spec.template.spec.initContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

# .spec.template.spec.initContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.initContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.initContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.initContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.initContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.initContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.initContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.initContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.initContainers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.initContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.initContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.template.spec.initContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.initContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.initContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.initContainers[].securityContext.appArmorProfile

### Description

AppArmorProfile defines a pod or container's AppArmor settings.

### Type

`object`

### Required

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.initContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.initContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.initContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.initContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.initContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.initContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.initContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.nodeSelector

**Description**

NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

**Type**

`object`

## .spec.template.spec.os

**Description**

PodOS defines the OS parameters of a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration ↗ Clients should expect to handle additional values and treat unrecognized values in this field as os: null |

## .spec.template.spec.overhead

**Description**

Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md

**Type**

`object`

## .spec.template.spec.readinessGates

**Description**

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates

**Type**

`array`

## .spec.template.spec.readinessGates[]

**Description**

PodReadinessGate contains the reference to a pod condition

**Type**

`object`

**Required**

`conditionType`

| Property | Type | Description |
|---|---|---|
| `conditionType` | `string` | ConditionType refers to a condition in the pod's condition list with matching type. |

## .spec.template.spec.resourceClaims

**Description**

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable.

**Type**

`array`

## .spec.template.spec.resourceClaims[]

**Description**

PodResourceClaim references exactly one ResourceClaim, either directly or by naming a ResourceClaimTemplate which is then turned into a ResourceClaim for the pod. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL. |
| `resourceClaimName` | `string` | ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |
| `resourceClaimTemplateName` | `string` | ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.<br><br>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The pod name and resource name, along with a generated component, will be used to form a unique name for the ResourceClaim, which will be recorded in pod.status.resourceClaimStatuses.<br><br>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |

## .spec.template.spec.resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.template.spec.resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |

| Property | Type | Description |
| --- | --- | --- |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.schedulingGates

**Description**

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod. SchedulingGates can only be set at pod creation time, and be removed only afterwards.

**Type**

`array`

## .spec.template.spec.schedulingGates[]

**Description**

PodSchedulingGate is associated to a Pod to guard its scheduling.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the scheduling gate. Each scheduling gate must have a unique name field. |

## .spec.template.spec.securityContext

**Description**

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `fsGroup` | `integer` | A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:<br>1. The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw----<br>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows. |
| `fsGroupChangePolicy` | `string` | fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership(and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.<br>Possible enum values:<br>- `"Always"` indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior.<br>- `"OnRootMismatch"` indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `seLinuxChangePolicy` | `string` | seLinuxChangePolicy defines how the container's SELinux label is applied to all volumes used by the Pod. It has no effect on nodes that do not support SELinux or to volumes does not support SELinux. Valid values are "MountOption" and "Recursive".<br><br>"Recursive" means relabeling of all files on all Pod volumes by the container runtime. This may be slow for large volumes, but allows mixing privileged and unprivileged Pods sharing the same volume on the same node.<br><br>"MountOption" mounts all eligible Pod volumes with `-o context` mount option. This requires all Pods that share the same volume to use the same SELinux label. It is not possible to share the same volume among privileged and unprivileged Pods. Eligible volumes are in-tree FibreChannel and iSCSI volumes, and all CSI volumes whose CSI driver announces SELinux support by setting spec.seLinuxMount: true in their CSIDriver instance. Other volumes are always re-labelled recursively. "MountOption" value is allowed only when SELinuxMount feature gate is enabled.<br><br>If not specified and SELinuxMount feature gate is enabled, "MountOption" is used. If not specified and SELinuxMount feature gate is disabled, "MountOption" is used for ReadWriteOncePod volumes and "Recursive" for all other volumes.<br><br>This field affects only Pods that have SELinux label set, either in PodSecurityContext or in SecurityContext of all containers.<br><br>All Pods that use the same volume should use the same seLinuxChangePolicy, otherwise some pods can get stuck in ContainerCreating state. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `supplementalGroups` | `array` | A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows. |
| `supplementalGroupsPolicy` | `string` | Defines how supplemental groups of the first container processes are calculated. Valid values are "Merge" and "Strict". If not specified, "Merge" is used. (Alpha) Using the field requires the SupplementalGroupsPolicy feature gate to be enabled and the container runtime must implement support for this feature. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br><br>- `"Merge"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be merged with the primary user's groups as defined in the container image (in /etc/group). |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Strict"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be used instead of any groups defined in the container image. |
| `sysctls` | `array` | Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br>• `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>• `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `level` | `string` | Level is SELinux level label that applies to the container. |

| Property | Type | Description |
| --- | --- | --- |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.securityContext.supplementalGroups

**Description**

A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.supplementalGroups[]

**Type**

`integer`

## .spec.template.spec.securityContext.sysctls

**Description**

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.sysctls[]

**Description**

Sysctl defines a kernel parameter to be set

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of a property to set |
| `value` | `string` | Value of a property to set |

## .spec.template.spec.securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |

| Property | Type | Description |
| --- | --- | --- |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.tolerations

**Description**

If specified, the pod's tolerations.

**Type**

`array`

## .spec.template.spec.tolerations[]

**Description**

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `effect` | `string` | Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.<br><br>Possible enum values:<br>- `"NoExecute"` Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController.<br>- `"NoSchedule"` Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler.<br>- `"PreferNoSchedule"` Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler. |
| `key` | `string` | Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys. |
| `operator` | `string` | Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"Equal"`<br>- `"Exists"` |
| `tolerationSeconds` | `integer` | TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system. |
| `value` | `string` | Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string. |

## .spec.template.spec.topologySpreadConstraints

**Description**

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[]

**Description**

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

**Type**

`object`

**Required**

`maxSkew` `topologyKey` `whenUnsatisfiable`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.<br><br>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `maxSkew` | `integer` | MaxSkew describes the degree to which pods may be unevenly distributed. When `whenUnsatisfiable=DoNotSchedule`, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1. \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P \| - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When `whenUnsatisfiable=ScheduleAnyway`, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed. |
| `minDomains` | `integer` | MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.<br><br>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2: \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P P \| The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew. |
| `nodeAffinityPolicy` | `string` | NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.<br><br>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |
| `nodeTaintsPolicy` | `string` | NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.<br><br>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |

| Property | Type | Description |
| --- | --- | --- |
| `topologyKey` | `string` | TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field. |
| `whenUnsatisfiable` | `string` | WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: \| zone1 \| zone2 \| zone3 \| \| P P P \| P \| P \| If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it *more* imbalanced. It's a required field.<br><br>Possible enum values:<br>- `"DoNotSchedule"` instructs the scheduler not to schedule the pod when constraints are not satisfied.<br>- `"ScheduleAnyway"` instructs the scheduler to schedule the pod even if constraints are not satisfied. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector. This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.volumes

**Description**

List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

**Type**

`array`

## .spec.template.spec.volumes[]

**Description**

Volume represents a named volume in a pod that may be accessed by any container in the pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `awsElasticBlockStore` | `object` | Represents a Persistent Disk resource in AWS.<br><br>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling. |
| `azureDisk` | `object` | AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod. |
| `azureFile` | `object` | AzureFile represents an Azure File Service mount on the host and bind mount to the pod. |
| `cephfs` | `object` | Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling. |
| `cinder` | `object` | Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership |

| Property | Type | Description |
| --- | --- | --- |
| | | management and SELinux relabeling. |
| `configMap` | `object` | Adapts a ConfigMap into a volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling. |
| `csi` | `object` | Represents a source location of a volume to mount, managed by an external CSI driver |
| `downwardAPI` | `object` | DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling. |
| `emptyDir` | `object` | Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling. |
| `ephemeral` | `object` | Represents an ephemeral volume that is handled by a normal storage driver. |
| `fc` | `object` | Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling. |
| `flexVolume` | `object` | FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin. |
| `flocker` | `object` | Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling. |
| `gcePersistentDisk` | `object` | Represents a Persistent Disk resource in Google Compute Engine.<br><br>A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling. |
| `gitRepo` | `object` | Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.<br><br>DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container. |

| Property | Type | Description |
| --- | --- | --- |
| `glusterfs` | `object` | Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling. |
| `hostPath` | `object` | Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling. |
| `image` | `object` | ImageVolumeSource represents a image volume resource. |
| `iscsi` | `object` | Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling. |
| `name` | `string` | name of the volume. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `nfs` | `object` | Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling. |
| `persistentVolumeClaim` | `object` | PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system). |
| `photonPersistentDisk` | `object` | Represents a Photon Controller persistent disk resource. |
| `portworxVolume` | `object` | PortworxVolumeSource represents a Portworx volume resource. |
| `projected` | `object` | Represents a projected volume source |
| `quobyte` | `object` | Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling. |
| `rbd` | `object` | Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling. |
| `scaleIO` | `object` | ScaleIOVolumeSource represents a persistent ScaleIO volume |

| Property | Type | Description |
| --- | --- | --- |
| `secret` | `object` | Adapts a Secret into a volume.<br><br>The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling. |
| `storageos` | `object` | Represents a StorageOS persistent volume resource. |
| `vsphereVolume` | `object` | Represents a vSphere volume resource. |

## .spec.template.spec.volumes[].awsElasticBlockStore

**Description**

Represents a Persistent Disk resource in AWS. An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). |
| `readOnly` | `boolean` | readOnly value true will force the readOnly setting in VolumeMounts. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `volumeID` | `string` | volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |

## .spec.template.spec.volumes[].azureDisk

**Description**

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`diskName` `diskURI`

| Property | Type | Description |
| --- | --- | --- |
| `cachingMode` | `string` | cachingMode is the Host Caching mode: None, Read Only, Read Write.<br><br>Possible enum values:<br>- `"None"`<br>- `"ReadOnly"`<br>- `"ReadWrite"` |
| `diskName` | `string` | diskName is the Name of the data disk in the blob storage |
| `diskURI` | `string` | diskURI is the URI of data disk in the blob storage |
| `fsType` | `string` | fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `kind` | `string` | kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared<br><br>Possible enum values:<br>- `"Dedicated"`<br>- `"Managed"`<br>- `"Shared"` |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

## .spec.template.spec.volumes[].azureFile

**Description**

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`secretName` `shareName`

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretName` | `string` | secretName is the name of secret that contains Azure Storage Account Name and Key |
| `shareName` | `string` | shareName is the azure share Name |

## .spec.template.spec.volumes[].cephfs

### Description

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`monitors`

| Property | Type | Description |
| --- | --- | --- |
| `monitors` | `array` | monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `path` | `string` | path is Optional: Used as the mounted root, rather than the full Ceph tree, default is / |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretFile` | `string` | secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].cephfs.monitors

### Description

monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].cephfs.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].cephfs.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].cinder

**Description**

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeID` | `string` | volumeID used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |

## .spec.template.spec.volumes[].cinder.secretRef

### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].configMap

### Description

Adapts a ConfigMap into a volume. The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].configMap.items

### Description

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].csi

**Description**

Represents a source location of a volume to mount, managed by an external CSI driver

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
|---|---|---|
| `driver` | `string` | driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster. |
| `fsType` | `string` | fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply. |

| Property | Type | Description |
| --- | --- | --- |
| `nodePublishSecretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `readOnly` | `boolean` | readOnly specifies a read-only configuration for the volume. Defaults to false (read/write). |
| `volumeAttributes` | `object` | volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values. |

## .spec.template.spec.volumes[].csi.nodePublishSecretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].csi.volumeAttributes

**Description**

volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**Type**

`object`

## .spec.template.spec.volumes[].downwardAPI

**Description**

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | Items is a list of downward API volume file |

## .spec.template.spec.volumes[].downwardAPI.items

**Description**

Items is a list of downward API volume file

**Type**

`array`

## .spec.template.spec.volumes[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
|---|---|---|
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].emptyDir

### Description

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `medium` | `string` | medium represents what type of storage medium should back this directory. The default is "" which means to use the n |
| `sizeLimit` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAM<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits> <r<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.html)<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a number<br>When a Quantity is parsed from a string, it will remember the type of suffix it had, an<br>Before serializing, Quantity will be put in "canonical form". This means that Exponent/<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suffix<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point number.<br>Non-canonical values will still parse as long as they are well formed, but will be re-e<br>This format is intended to make it difficult to use these numbers without writing some |

## .spec.template.spec.volumes[].ephemeral

**Description**

Represents an ephemeral volume that is handled by a normal storage driver.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `volumeClaimTemplate` | `object` | PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate

**Description**

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

**Type**

`object`

**Required**

`spec`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec

### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `accessModes` | `array` | accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1 ↗ |
| `dataSource` | `object` | TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace. |
| `dataSourceRef` | `object` | TypedObjectReference contains enough information to let you locate the typed referenced object |
| `resources` | `object` | VolumeResourceRequirements describes the storage resource requirements for a volume. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `storageClassName` | `string` | storageClassName is the name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1 ↗ |
| `volumeAttributesClassName` | `string` | volumeAttributesClassName may be used to set the VolumeAttributesClass used by this claim. If specified, the CSI driver will create or update the volume with the attributes defined in the corresponding VolumeAttributesClass. This has a different purpose than storageClassName, it can be changed after the claim is created. An empty string value means that no VolumeAttributesClass will be applied to the claim but it's not allowed to reset this field to empty string once it is set. If unspecified and the PersistentVolumeClaim is unbound, the default VolumeAttributesClass will be set by the persistentvolume controller if it exists. If the resource referred to by volumeAttributesClass does not exist, this PersistentVolumeClaim will be set to a Pending state, as reflected by the |

| Property | Type | Description |
| --- | --- | --- |
| | | modifyVolumeStatus field, until such as a resource exists. More info: https://kubernetes.io/docs/concepts/storage/volume-attributes-classes/ (Beta) Using this field requires the VolumeAttributesClass feature gate to be enabled (off by default). |
| `volumeMode` | `string` | volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.<br><br>Possible enum values:<br>• `"Block"` means the volume will not be formatted with a filesystem and will remain a raw block device.<br>• `"Filesystem"` means the volume will be or is formatted with a filesystem. |
| `volumeName` | `string` | volumeName is the binding reference to the PersistentVolume backing this claim. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes

**Description**

accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

**Description**

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is the name of resource being referenced |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef

### Description

TypedObjectReference contains enough information to let you locate the typed referenced object

### Type

`object`

### Required

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |
| `namespace` | `string` | Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

### Description

VolumeResourceRequirements describes the storage resource requirements for a volume.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].fc

**Description**

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `lun` | `integer` | lun is Optional: FC target lun number |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `targetWWNs` | `array` | targetWWNs is Optional: FC target worldwide names (WWNs) |
| `wwids` | `array` | wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously. |

## .spec.template.spec.volumes[].fc.targetWWNs

**Description**

targetWWNs is Optional: FC target worldwide names (WWNs)

**Type**

`array`

## .spec.template.spec.volumes[].fc.targetWWNs[]

**Type**

`string`

## .spec.template.spec.volumes[].fc.wwids

**Description**

wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

**Type**

`array`

## .spec.template.spec.volumes[].fc.wwids[]

**Type**

`string`

## .spec.template.spec.volumes[].flexVolume

**Description**

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
| --- | --- | --- |
| `driver` | `string` | driver is the name of the driver to use for this volume. |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script. |
| `options` | `object` | options is Optional: this field holds extra command options if any. |
| `readOnly` | `boolean` | readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |

## .spec.template.spec.volumes[].flexVolume.options

**Description**

options is Optional: this field holds extra command options if any.

**Type**

`object`

## .spec.template.spec.volumes[].flexVolume.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: |

| Property | Type | Description |
| --- | --- | --- |
| | | https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].flocker

### Description

Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `datasetName` | `string` | datasetName is Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated |
| `datasetUUID` | `string` | datasetUUID is the UUID of the dataset. This is unique identifier of a Flocker dataset |

## .spec.template.spec.volumes[].gcePersistentDisk

### Description

Represents a Persistent Disk resource in Google Compute Engine. A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

### Type

`object`

### Required

`pdName`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `pdName` | `string` | pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk |

## .spec.template.spec.volumes[].gitRepo

**Description**

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

**Type**

`object`

**Required**

`repository`

| Property | Type | Description |
| --- | --- | --- |
| `directory` | `string` | directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name. |
| `repository` | `string` | repository is the URL |
| `revision` | `string` | revision is the commit hash for the specified revision. |

## .spec.template.spec.volumes[].glusterfs

**Description**

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`endpoints` `path`

| Property | Type | Description |
| --- | --- | --- |
| `endpoints` | `string` | endpoints is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod |
| `path` | `string` | path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

## .spec.template.spec.volumes[].hostPath

### Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`path`

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗ |
| `type` | `string` | type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗<br><br>Possible enum values:<br>• `""` For backwards compatible, leave it empty if unset<br>• `"BlockDevice"` A block device must exist at the given path<br>• `"CharDevice"` A character device must exist at the given path<br>• `"Directory"` A directory must exist at the given path<br>• `"DirectoryOrCreate"` If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet.<br>• `"File"` A file must exist at the given path<br>• `"FileOrCreate"` If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet.<br>• `"Socket"` A UNIX socket must exist at the given path |

## .spec.template.spec.volumes[].image

### Description

ImageVolumeSource represents a image volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `pullPolicy` | `string` | Policy for pulling OCI objects. Possible values are: Always: the kubelet always attempts to pull the reference. Container creation will fail If the pull fails. Never: the kubelet never pulls the reference and only uses a local image or artifact. |

| Property | Type | Description |
| --- | --- | --- |
| | | Container creation will fail if the reference isn't present. IfNotPresent: the kubelet pulls if the reference isn't already present on disk. Container creation will fail if the reference isn't present and the pull fails. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise.<br><br>Possible enum values:<br><br>- `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>- `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>- `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `reference` | `string` | Required: Image or artifact reference to be used. Behaves in the same way as pod.spec.containers[*].image. Pull secrets will be assembled in the same way as for the container image by looking up node credentials, SA image pull secrets, and pod spec image pull secrets. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |

## .spec.template.spec.volumes[].iscsi

### Description

Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`targetPortal` `iqn` `lun`

| Property | Type | Description |
| --- | --- | --- |
| `chapAuthDiscovery` | `boolean` | chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication |
| `chapAuthSession` | `boolean` | chapAuthSession defines whether support iSCSI Session CHAP authentication |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi ↗ |
| `initiatorName` | `string` | initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection. |
| `iqn` | `string` | iqn is the target iSCSI Qualified Name. |

| Property | Type | Description |
| --- | --- | --- |
| `iscsiInterface` | `string` | iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp). |
| `lun` | `integer` | lun represents iSCSI Target Lun number. |
| `portals` | `array` | portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `targetPortal` | `string` | targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |

## .spec.template.spec.volumes[].iscsi.portals

**Description**

portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

**Type**

`array`

## .spec.template.spec.volumes[].iscsi.portals[]

**Type**

`string`

## .spec.template.spec.volumes[].iscsi.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].nfs

**Description**

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`server` `path`

| Property | Type | Description |
|---|---|---|
| `path` | `string` | path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |
| `readOnly` | `boolean` | readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |
| `server` | `string` | server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |

## .spec.template.spec.volumes[].persistentVolumeClaim

**Description**

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

**Type**

`object`

**Required**

`claimName`

| Property | Type | Description |
|---|---|---|
| `claimName` | `string` | claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims |
| `readOnly` | `boolean` | readOnly Will force the ReadOnly setting in VolumeMounts. Default false. |

## .spec.template.spec.volumes[].photonPersistentDisk

**Description**

Represents a Photon Controller persistent disk resource.

**Type**

`object`

**Required**

`pdID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `pdID` | `string` | pdID is the ID that identifies Photon Controller persistent disk |

## .spec.template.spec.volumes[].portworxVolume

**Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `volumeID` | `string` | volumeID uniquely identifies a Portworx volume |

## .spec.template.spec.volumes[].projected

**Description**

Represents a projected volume source

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `defaultMode` | `integer` | defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `sources` | `array` | sources is the list of volume projections. Each entry in this list handles one source. |

## .spec.template.spec.volumes[].projected.sources

**Description**

sources is the list of volume projections. Each entry in this list handles one source.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[]

**Description**

Projection that may be projected along with other supported volume types. Exactly one of these fields must be set.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `clusterTrustBundle` | `object` | ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem. |
| `configMap` | `object` | Adapts a ConfigMap into a projected volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode. |
| `downwardAPI` | `object` | Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode. |
| `secret` | `object` | Adapts a secret into a projected volume.<br><br>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode. |
| `serviceAccountToken` | `object` | ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise). |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle

**Description**

ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem.

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `name` | `string` | Select a single ClusterTrustBundle by object name. Mutually-exclusive with signerName and labelSelector. |
| `optional` | `boolean` | If true, don't block pod startup if the referenced ClusterTrustBundle(s) aren't available. If using name, then the named ClusterTrustBundle is allowed not to exist. If using signerName, then the combination of signerName and labelSelector is allowed to match zero ClusterTrustBundles. |
| `path` | `string` | Relative path from the volume root to write the bundle. |
| `signerName` | `string` | Select all ClusterTrustBundles that match this signer name. Mutually-exclusive with name. The contents of all selected ClusterTrustBundles will be unified and deduplicated. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].projected.sources[].configMap

**Description**

Adapts a ConfigMap into a projected volume. The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].projected.sources[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI

**Description**

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | Items is a list of DownwardAPIVolume file |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items

**Description**

Items is a list of DownwardAPIVolume file

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |

| Property | Type | Description |
| --- | --- | --- |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and |

| Property | Type | Description |
| --- | --- | --- |
| | | The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].projected.sources[].secret

### Description

Adapts a secret into a projected volume. The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | optional field specify whether the Secret or its key must be defined |

## .spec.template.spec.volumes[].projected.sources[].secret.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].serviceAccountToken

**Description**

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `audience` | `string` | audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver. |
| `expirationSeconds` | `integer` | expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours.Defaults to 1 hour and must be at least 10 minutes. |
| `path` | `string` | path is the path relative to the mount point of the file to project the token into. |

## .spec.template.spec.volumes[].quobyte

**Description**

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`registry` `volume`

| Property | Type | Description |
|---|---|---|
| `group` | `string` | group to map volume access to Default is no group |
| `readOnly` | `boolean` | readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false. |
| `registry` | `string` | registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes |
| `tenant` | `string` | tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin |
| `user` | `string` | user to map volume access to Defaults to serivceaccount user |

| Property | Type | Description |
| --- | --- | --- |
| `volume` | `string` | volume is a string that references an already created Quobyte volume by name. |

## .spec.template.spec.volumes[].rbd

### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`monitors` `image`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd ↗ |
| `image` | `string` | image is the rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `keyring` | `string` | keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `monitors` | `array` | monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `pool` | `string` | pool is the rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is the rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].rbd.monitors

**Description**

monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].rbd.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].rbd.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].scaleIO

**Description**

ScaleIOVolumeSource represents a persistent ScaleIO volume

**Type**

`object`

**Required**

`gateway` `system` `secretRef`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs". |
| `gateway` | `string` | gateway is the host address of the ScaleIO API Gateway. |
| `protectionDomain` | `string` | protectionDomain is the name of the ScaleIO Protection Domain for the configured storage. |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

| Property | Type | Description |
| --- | --- | --- |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `sslEnabled` | `boolean` | sslEnabled Flag enable/disable SSL communication with Gateway, default false |
| `storageMode` | `string` | storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned. |
| `storagePool` | `string` | storagePool is the ScaleIO Storage Pool associated with the protection domain. |
| `system` | `string` | system is the name of the storage system as configured in ScaleIO. |
| `volumeName` | `string` | volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source. |

## .spec.template.spec.volumes[].scaleIO.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].secret

**Description**

Adapts a Secret into a volume. The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON |

| Property | Type | Description |
| --- | --- | --- |
| | | requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `optional` | `boolean` | optional field specify whether the Secret or its keys must be defined |
| `secretName` | `string` | secretName is the name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret ↗ |

## .spec.template.spec.volumes[].secret.items

**Description**

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].storageos

### Description

Represents a StorageOS persistent volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeName` | `string` | volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace. |
| `volumeNamespace` | `string` | volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created. |

## .spec.template.spec.volumes[].storageos.secretRef

### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].vsphereVolume

**Description**

Represents a vSphere volume resource.

**Type**

`object`

**Required**

`volumePath`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `storagePolicyID` | `string` | storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName. |
| `storagePolicyName` | `string` | storagePolicyName is the storage Policy Based Management (SPBM) profile name. |
| `volumePath` | `string` | volumePath is the path that identifies vSphere volume vmdk |

## .status

**Description**

JobStatus represents the current state of a Job.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `active` | `integer` | The number of pending and running pods which are not terminating (without a deletionTimestamp). The value is zero for finished jobs. |
| `completedIndexes` | `string` | completedIndexes holds the completed indexes when .spec.completionMode = "Indexed" in a text format. The indexes are represented as decimal integers separated by commas. The numbers are listed in increasing order. Three or more consecutive numbers are compressed and represented by the first and last element of the series, separated by a hyphen. For example, if the completed indexes are 1, 3, 4, 5 and 7, they are represented as "1,3-5,7". |
| `completionTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

| Property | Type | Description |
| --- | --- | --- |
| `conditions` | `array` | The latest available observations of an object's current state. When a Job fails, one of the conditions will have type "Failed" and status true. When a Job is suspended, one of the conditions will have type "Suspended" and status true; when the Job is resumed, the status of this condition will become false. When a Job is completed, one of the conditions will have type "Complete" and status true.<br><br>A job is considered finished when it is in a terminal condition, either "Complete" or "Failed". A Job cannot have both the "Complete" and "Failed" conditions. Additionally, it cannot be in the "Complete" and "FailureTarget" conditions. The "Complete", "Failed" and "FailureTarget" conditions cannot be disabled.<br><br>More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/ ↗ |
| `failed` | `integer` | The number of pods which reached phase Failed. The value increases monotonically. |
| `failedIndexes` | `string` | FailedIndexes holds the failed indexes when spec.backoffLimitPerIndex is set. The indexes are represented in the text format analogous as for the `completedIndexes` field, ie. they are kept as decimal integers separated by commas. The numbers are listed in increasing order. Three or more consecutive numbers are compressed and represented by the first and last element of the series, separated by a hyphen. For example, if the failed indexes are 1, 3, 4, 5 and 7, they are represented as "1,3-5,7". The set of failed indexes cannot overlap with the set of completed indexes.<br><br>This field is beta-level. It can be used when the `JobBackoffLimitPerIndex` feature gate is enabled (enabled by default). |
| `ready` | `integer` | The number of active pods which have a Ready condition and are not terminating (without a deletionTimestamp). |
| `startTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `succeeded` | `integer` | The number of pods which reached phase Succeeded. The value increases monotonically for a given spec. However, it may decrease in reaction to scale down of elastic indexed jobs. |
| `terminating` | `integer` | The number of pods which are terminating (in phase Pending or Running and have a deletionTimestamp).<br><br>This field is beta-level. The job controller populates the field when the feature gate JobPodReplacementPolicy is enabled (enabled by default). |
| `uncountedTerminatedPods` | `object` | UncountedTerminatedPods holds UIDs of Pods that have terminated but haven't been accounted in Job status counters. |

## .status.conditions

**Description**

The latest available observations of an object's current state. When a Job fails, one of the conditions will have type "Failed" and status true. When a Job is suspended, one of the conditions will have type "Suspended" and status true; when the Job is resumed, the status of this condition will become false. When a Job is completed, one of the conditions will have type "Complete" and status true. A job is considered finished when it is in a terminal condition, either "Complete" or "Failed". A Job cannot have both the "Complete" and "Failed" conditions. Additionally, it cannot be in the "Complete" and "FailureTarget" conditions. The "Complete", "Failed" and "FailureTarget" conditions cannot be disabled. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

**Type**

`array`

## .status.conditions[]

**Description**

JobCondition describes current state of a job.

**Type**

`object`

**Required**

`type` `status`

| Property | Type | Description |
| --- | --- | --- |
| `lastProbeTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `lastTransitionTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | Human readable message indicating details about last transition. |
| `reason` | `string` | (brief) reason for the condition's last transition. |
| `status` | `string` | Status of the condition, one of True, False, Unknown. |
| `type` | `string` | Type of job condition, Complete or Failed. |

## .status.uncountedTerminatedPods

**Description**

UncountedTerminatedPods holds UIDs of Pods that have terminated but haven't been accounted in Job status counters.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `failed` | `array` | failed holds UIDs of failed Pods. |
| `succeeded` | `array` | succeeded holds UIDs of succeeded Pods. |

### .status.uncountedTerminatedPods.failed

**Description**

failed holds UIDs of failed Pods.

**Type**

`array`

### .status.uncountedTerminatedPods.failed[]

**Type**

`string`

### .status.uncountedTerminatedPods.succeeded

**Description**

succeeded holds UIDs of succeeded Pods.

**Type**

`array`

### .status.uncountedTerminatedPods.succeeded[]

**Type**

`string`

## API Endpoints

The following API endpoints are available:

- `/kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/jobs`
  - `DELETE`: delete collection of Job
  - `GET`: list objects of kind Job
  - `POST`: create a new Job
- `/kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/jobs/{name}`
  - `DELETE`: delete the specified Job
  - `GET`: read the specified Job
  - `PATCH`: partially update the specified Job
  - `PUT`: replace the specified Job

- `/kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/jobs/{name}/status`
  - `GET`: read status of the specified Job
  - `PATCH`: partially update status of the specified Job
  - `PUT`: replace status of the specified Job

## /kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/jobs

### HTTP method

`DELETE`

### Description

delete collection of Job

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

### Description

list objects of kind Job

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `JobList` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`POST`

### Description

create a new Job

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `Job` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Job` schema |
| 201 - Created | `Job` schema |
| 202 - Accepted | `Job` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/jobs/{name}

### HTTP method

`DELETE`

### Description

delete the specified Job

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 202 - Accepted | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

### Description

read the specified Job

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Job` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update the specified Job

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `Job` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace the specified Job

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
|---|---|---|
| `body` | `Job` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Job` schema |
| 201 - Created | `Job` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/batch/v1/namespaces/{namespace}/jobs/{name}/status

### HTTP method

`GET`

### Description

read status of the specified Job

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Job` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update status of the specified Job

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Job` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace status of the specified Job

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `Job` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Job` schema |
| 201 - Created | `Job` schema |
| 401 - Unauthorized | Empty |

# Pod [v1]

**Description**

Pod is a collection of containers that can run on a host. This resource is created by clients and scheduled onto hosts.

**Type**

`object`

## Specification

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources |
| `kind` | `string` | Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PodSpec is a description of a pod. |
| `status` | `object` | PodStatus represents information about the status of a pod. Status may trail the actual state of a system, especially if the node that hosts the pod cannot contact the control plane. |

### .spec

**Description**

PodSpec is a description of a pod.

**Type**

`object`

**Required**

`containers`

| Property | Type | Description |
| --- | --- | --- |
| `activeDeadlineSeconds` | `integer` | Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer. |
| `affinity` | `object` | Affinity is a group of affinity scheduling rules. |
| `automountServiceAccountToken` | `boolean` | AutomountServiceAccountToken indicates whether a service account token should be automatically mounted. |
| `containers` | `array` | List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated. |
| `dnsConfig` | `object` | PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy. |
| `dnsPolicy` | `string` | Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.<br><br>Possible enum values:<br>• `"ClusterFirst"` indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>• `"ClusterFirstWithHostNet"` indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>• `"Default"` indicates that the pod should use the default (as determined by kubelet) DNS settings.<br>• `"None"` indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig. |
| `enableServiceLinks` | `boolean` | EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching traditional container linking syntax. Optional: Defaults to true. |
| `ephemeralContainers` | `array` | List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. |

| Property | Type | Description |
| --- | --- | --- |
| `hostAliases` | `array` | HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. |
| `hostIPC` | `boolean` | Use the host's ipc namespace. Optional: Default to false. |
| `hostNetwork` | `boolean` | Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false. |
| `hostPID` | `boolean` | Use the host's pid namespace. Optional: Default to false. |
| `hostUsers` | `boolean` | Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new userns is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature. |
| `hostname` | `string` | Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value. |
| `imagePullSecrets` | `array` | ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod ↗ |
| `initContainers` | `array` | List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/ ↗ |
| `nodeName` | `string` | NodeName indicates in which node this pod is scheduled. If empty, this pod is a candidate for scheduling by the scheduler defined in schedulerName. Once this field is set, the kubelet for this node becomes responsible for the lifecycle of this pod. This field should not be used to express a desire for the pod to be scheduled on a specific node. https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodename ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `nodeSelector` | `object` | NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/ |
| `os` | `object` | PodOS defines the OS parameters of a pod. |
| `overhead` | `object` | Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md |
| `preemptionPolicy` | `string` | PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.<br><br>Possible enum values:<br>• `"Never"` means that pod never preempts other pods with lower priority.<br>• `"PreemptLowerPriority"` means that pod can preempt other pods with lower priority. |
| `priority` | `integer` | The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority. |
| `priorityClassName` | `string` | If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default. |
| `readinessGates` | `array` | If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates |

| Property | Type | Description |
| --- | --- | --- |
| `resourceClaims` | `array` | ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.<br><br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy ↗<br><br>Possible enum values:<br>• `"Always"`<br>• `"Never"`<br>• `"OnFailure"` |
| `runtimeClassName` | `string` | RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class ↗ |
| `schedulerName` | `string` | If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler. |
| `schedulingGates` | `array` | SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.<br><br>SchedulingGates can only be set at pod creation time, and be removed only afterwards. |
| `securityContext` | `object` | PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext. |
| `serviceAccount` | `string` | DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead. |

| Property | Type | Description |
| --- | --- | --- |
| `serviceAccountName` | `string` | ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/ |
| `setHostnameAsFQDN` | `boolean` | If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false. |
| `shareProcessNamespace` | `boolean` | Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false. |
| `subdomain` | `string` | If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all. |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds. |
| `tolerations` | `array` | If specified, the pod's tolerations. |
| `topologySpreadConstraints` | `array` | TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed. |
| `volumes` | `array` | List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes |

## .spec.affinity

### Description

Affinity is a group of affinity scheduling rules.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `nodeAffinity` | `object` | Node affinity is a group of node affinity scheduling rules. |
| `podAffinity` | `object` | Pod affinity is a group of inter pod affinity scheduling rules. |
| `podAntiAffinity` | `object` | Pod anti affinity is a group of inter pod anti affinity scheduling rules. |

## .spec.affinity.nodeAffinity

**Description**

Node affinity is a group of node affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `object` | A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms. |

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

**Type**

`object`

**Required**

`weight` `preference`

| Property | Type | Description |
|---|---|---|
| `preference` | `object` | A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm. |
| `weight` | `integer` | Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100. |

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br><br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values[]

**Type**

`string`

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]

**Type**

`string`

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

**Type**

`object`

**Required**

`nodeSelectorTerms`

| Property | Type | Description |
|---|---|---|
| `nodeSelectorTerms` | `array` | Required. A list of node selector terms. The terms are ORed. |

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms

**Description**

Required. A list of node selector terms. The terms are ORed.

**Type**

`array`

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]

**Type**

`string`

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values[]

**Type**

`string`

## .spec.affinity.podAffinity

**Description**

Pod affinity is a group of inter pod affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
| --- | --- | --- |
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey>

matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.affinity.podAntiAffinity

**Description**

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
|---|---|---|
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only |

| Property | Type | Description |
| --- | --- | --- |
| | | "value". The requirements are ANDed. |

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |

| Property | Type | Description |
| --- | --- | --- |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.containers

### Description

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

### Type

`array`

## .spec.containers[]

### Description

A single application container that you want to run within a pod.

### Type

`object`

### Required

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |

| Property | Type | Description |
| --- | --- | --- |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |

| Property | Type | Description |
| --- | --- | --- |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values:<br><br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.containers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.containers[].args[]

**Type**

`string`

## .spec.containers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.containers[].command[]

**Type**

`string`

## .spec.containers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.containers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.containers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

| Property | Type | Description |
| --- | --- | --- |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.containers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.containers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.containers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.containers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.containers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.containers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.containers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.containers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.containers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

| Property | Type | Description |
| --- | --- | --- |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.containers[].lifecycle.postStart

### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.containers[].lifecycle.postStart.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.containers[].lifecycle.postStart.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.containers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.containers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.containers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.containers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
|---|---|---|
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.containers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.containers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.containers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.containers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to

explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.containers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.containers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.containers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.containers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
|---|---|---|
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.containers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.containers[].livenessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling |

| Property | Type | Description |
| --- | --- | --- |
| | | ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.containers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.containers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.containers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.containers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.containers[].livenessProbe.httpGet

### Description

HTTPGetAction describes an action based on HTTP Get requests.

### Type

`object`

### Required

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.containers[].livenessProbe.httpGet.httpHeaders

### Description

Custom headers to set in the request. HTTP allows repeated headers.

### Type

`array`

## .spec.containers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.containers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.containers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.containers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
|---|---|---|
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br><br>Possible enum values:<br>- `"SCTP"` is the SCTP protocol.<br>- `"TCP"` is the TCP protocol.<br>- `"UDP"` is the UDP protocol. |

## .spec.containers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

## .spec.containers[].readinessProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.containers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.containers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.containers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.containers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.containers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.containers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.containers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.containers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.containers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.containers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.containers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.containers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |

| Property | Type | Description |
| --- | --- | --- |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.containers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.containers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.containers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>- `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>- `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.containers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>- `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>- `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.containers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.containers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.containers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.containers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.containers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.containers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.containers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are: |

| Property | Type | Description |
| --- | --- | --- |
| | | Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>- `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>- `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.containers[].securityContext.windowsOptions

### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.containers[].startupProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.containers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.containers[].startupProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.containers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.containers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.containers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.containers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.containers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.containers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.containers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.containers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

# .spec.containers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

# .spec.containers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
|---|---|---|
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field |

| Property | Type | Description |
| --- | --- | --- |
| | | is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.dnsConfig

**Description**

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `nameservers` | `array` | A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed. |
| `options` | `array` | A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy. |
| `searches` | `array` | A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed. |

## .spec.dnsConfig.nameservers

**Description**

A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.

**Type**

`array`

## .spec.dnsConfig.nameservers[]

**Type**

`string`

## .spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

`array`

## .spec.dnsConfig.options[]

**Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is this DNS resolver option's name. Required. |
| `value` | `string` | Value is this DNS resolver option's value. |

## .spec.dnsConfig.searches

**Description**

A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

**Type**

`array`

## .spec.dnsConfig.searches[]

**Type**

`string`

## .spec.ephemeralContainers

**Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

`array`

## .spec.ephemeralContainers[]

**Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `args` | `array` | Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers. |
| `ports` | `array` | Ports are not allowed for ephemeral containers. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for the container to manage the restart behavior of each container within a pod. This may only be set for init containers. You cannot set this field on ephemeral containers. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to |

| Property | Type | Description |
| --- | --- | --- |
| | | stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `targetContainerName` | `string` | If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.<br><br>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined. |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values:<br><br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.ephemeralContainers[].args

**Description**

Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.ephemeralContainers[].args[]

**Type**

`string`

## .spec.ephemeralContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.ephemeralContainers[].command[]

**Type**

`string`

## .spec.ephemeralContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.ephemeralContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.ephemeralContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.ephemeralContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.ephemeralContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.ephemeralContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.ephemeralContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.ephemeralContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.ephemeralContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.ephemeralContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.ephemeralContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.ephemeralContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.ephemeralContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.ephemeralContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.ephemeralContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.ephemeralContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.ephemeralContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.ephemeralContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.ephemeralContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.ephemeralContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.ephemeralContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.ephemeralContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.ephemeralContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.ephemeralContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.ephemeralContainers[].livenessProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.ephemeralContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.ephemeralContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.ephemeralContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.ephemeralContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.ephemeralContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.ephemeralContainers[].ports

**Description**

Ports are not allowed for ephemeral containers.

**Type**

`array`

## .spec.ephemeralContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |

| Property | Type | Description |
| --- | --- | --- |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br><br>Possible enum values:<br>- `"SCTP"` is the SCTP protocol.<br>- `"TCP"` is the TCP protocol.<br>- `"UDP"` is the UDP protocol. |

## .spec.ephemeralContainers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.ephemeralContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.ephemeralContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.ephemeralContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.ephemeralContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.ephemeralContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.ephemeralContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.ephemeralContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.ephemeralContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.ephemeralContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.ephemeralContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.ephemeralContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.ephemeralContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.ephemeralContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.ephemeralContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.ephemeralContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.ephemeralContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.ephemeralContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.ephemeralContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.ephemeralContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.ephemeralContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.ephemeralContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.ephemeralContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.ephemeralContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.ephemeralContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.ephemeralContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.ephemeralContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.ephemeralContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.ephemeralContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.ephemeralContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.ephemeralContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.ephemeralContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.ephemeralContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.ephemeralContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

**Type**

`array`

## .spec.ephemeralContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.hostAliases

**Description**

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified.

**Type**

`array`

## .spec.hostAliases[]

**Description**

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

**Type**

`object`

**Required**

`ip`

| Property | Type | Description |
|---|---|---|
| `hostnames` | `array` | Hostnames for the above IP address. |
| `ip` | `string` | IP address of the host file entry. |

## .spec.hostAliases[].hostnames

**Description**

Hostnames for the above IP address.

**Type**

`array`

## .spec.hostAliases[].hostnames[]

**Type**

`string`

## .spec.imagePullSecrets

**Description**

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

**Type**

`array`

## .spec.imagePullSecrets[]

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.initContainers

### Description

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

### Type

`array`

## .spec.initContainers[]

### Description

A single application container that you want to run within a pod.

### Type

`object`

### Required

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |

| Property | Type | Description |
| --- | --- | --- |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.initContainers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.initContainers[].args[]

**Type**

`string`

## .spec.initContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.initContainers[].command[]

**Type**

`string`

## .spec.initContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.initContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.initContainers[].env[].valueFrom

### Description

EnvVarSource represents a source for the value of an EnvVar.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.initContainers[].env[].valueFrom.configMapKeyRef

### Description

Selects a key from a ConfigMap.

### Type

`object`

### Required

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.initContainers[].env[].valueFrom.fieldRef

### Description

ObjectFieldSelector selects an APIVersioned field of an object.

### Type

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.initContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.initContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.initContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.initContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.initContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.initContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.initContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.initContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.initContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.initContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.initContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.initContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.initContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.initContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.initContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.initContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.initContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.initContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.initContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.initContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.initContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.initContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.initContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.initContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.initContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.initContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.initContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.initContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.initContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.initContainers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.initContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.initContainers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.initContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.initContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.initContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.initContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.initContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.initContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.initContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.initContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.initContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.initContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.initContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.initContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.initContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.initContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.initContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>- `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>- `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.initContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.initContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.initContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.initContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.initContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.initContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.initContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.initContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.initContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.initContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.initContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.initContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.initContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.initContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.initContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.initContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.initContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.initContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.initContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.initContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.initContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.initContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.nodeSelector

**Description**

NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

**Type**

`object`

## .spec.os

**Description**

PodOS defines the OS parameters of a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration ↗ Clients should expect to handle additional values and treat unrecognized values in this field as os: null |

## .spec.overhead

**Description**

Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md

**Type**

`object`

## .spec.readinessGates

**Description**

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates

**Type**

`array`

## .spec.readinessGates[]

**Description**

PodReadinessGate contains the reference to a pod condition

**Type**

`object`

**Required**

`conditionType`

| Property | Type | Description |
|---|---|---|
| `conditionType` | `string` | ConditionType refers to a condition in the pod's condition list with matching type. |

## .spec.resourceClaims

**Description**

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable.

**Type**

`array`

## .spec.resourceClaims[]

**Description**

PodResourceClaim references exactly one ResourceClaim, either directly or by naming a ResourceClaimTemplate which is then turned into a ResourceClaim for the pod. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL. |
| `resourceClaimName` | `string` | ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |
| `resourceClaimTemplateName` | `string` | ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.<br><br>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The pod name and resource name, along with a generated component, will be used to form a unique name for the ResourceClaim, which will be recorded in pod.status.resourceClaimStatuses.<br><br>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |

## .spec.resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |

| Property | Type | Description |
| --- | --- | --- |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.schedulingGates

**Description**

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod. SchedulingGates can only be set at pod creation time, and be removed only afterwards.

**Type**

`array`

## .spec.schedulingGates[]

**Description**

PodSchedulingGate is associated to a Pod to guard its scheduling.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the scheduling gate. Each scheduling gate must have a unique name field. |

## .spec.securityContext

**Description**

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `fsGroup` | `integer` | A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:<br>1. The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw----<br>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows. |
| `fsGroupChangePolicy` | `string` | fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership(and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.<br>Possible enum values:<br>- `"Always"` indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior.<br>- `"OnRootMismatch"` indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `seLinuxChangePolicy` | `string` | seLinuxChangePolicy defines how the container's SELinux label is applied to all volumes used by the Pod. It has no effect on nodes that do not support SELinux or to volumes does not support SELinux. Valid values are "MountOption" and "Recursive".<br><br>"Recursive" means relabeling of all files on all Pod volumes by the container runtime. This may be slow for large volumes, but allows mixing privileged and unprivileged Pods sharing the same volume on the same node.<br><br>"MountOption" mounts all eligible Pod volumes with `-o context` mount option. This requires all Pods that share the same volume to use the same SELinux label. It is not possible to share the same volume among privileged and unprivileged Pods. Eligible volumes are in-tree FibreChannel and iSCSI volumes, and all CSI volumes whose CSI driver announces SELinux support by setting spec.seLinuxMount: true in their CSIDriver instance. Other volumes are always re-labelled recursively. "MountOption" value is allowed only when SELinuxMount feature gate is enabled.<br><br>If not specified and SELinuxMount feature gate is enabled, "MountOption" is used. If not specified and SELinuxMount feature gate is disabled, "MountOption" is used for ReadWriteOncePod volumes and "Recursive" for all other volumes.<br><br>This field affects only Pods that have SELinux label set, either in PodSecurityContext or in SecurityContext of all containers.<br><br>All Pods that use the same volume should use the same seLinuxChangePolicy, otherwise some pods can get stuck in ContainerCreating state. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `supplementalGroups` | `array` | A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows. |
| `supplementalGroupsPolicy` | `string` | Defines how supplemental groups of the first container processes are calculated. Valid values are "Merge" and "Strict". If not specified, "Merge" is used. (Alpha) Using the field requires the SupplementalGroupsPolicy feature gate to be enabled and the container runtime must implement support for this feature. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br><br>- `"Merge"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be merged with the primary user's groups as defined in the container image (in /etc/group). |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Strict"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be used instead of any groups defined in the container image. |
| `sysctls` | `array` | Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br>• `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>• `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `level` | `string` | Level is SELinux level label that applies to the container. |

| Property | Type | Description |
| --- | --- | --- |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.securityContext.supplementalGroups

**Description**

A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.securityContext.supplementalGroups[]

**Type**

`integer`

## .spec.securityContext.sysctls

**Description**

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.securityContext.sysctls[]

**Description**

Sysctl defines a kernel parameter to be set

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of a property to set |
| `value` | `string` | Value of a property to set |

## .spec.securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |

| Property | Type | Description |
| --- | --- | --- |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.tolerations

**Description**

If specified, the pod's tolerations.

**Type**

`array`

## .spec.tolerations[]

**Description**

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `effect` | `string` | Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.<br><br>Possible enum values:<br>- `"NoExecute"` Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController.<br>- `"NoSchedule"` Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler.<br>- `"PreferNoSchedule"` Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler. |
| `key` | `string` | Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys. |
| `operator` | `string` | Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"Equal"`<br>- `"Exists"` |
| `tolerationSeconds` | `integer` | TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system. |
| `value` | `string` | Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string. |

## .spec.topologySpreadConstraints

**Description**

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

**Type**

`array`

## .spec.topologySpreadConstraints[]

**Description**

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

**Type**

`object`

**Required**

`maxSkew` `topologyKey` `whenUnsatisfiable`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.<br><br>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `maxSkew` | `integer` | MaxSkew describes the degree to which pods may be unevenly distributed. When `whenUnsatisfiable=DoNotSchedule`, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1. \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P \| - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When `whenUnsatisfiable=ScheduleAnyway`, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed. |
| `minDomains` | `integer` | MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.<br><br>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2: \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P P \| The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew. |
| `nodeAffinityPolicy` | `string` | NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.<br><br>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br><br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |
| `nodeTaintsPolicy` | `string` | NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.<br><br>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br><br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |

| Property | Type | Description |
| --- | --- | --- |
| `topologyKey` | `string` | TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field. |
| `whenUnsatisfiable` | `string` | WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: \| zone1 \| zone2 \| zone3 \| \| P P P \| P \| P \| If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it *more* imbalanced. It's a required field.<br><br>Possible enum values:<br>- `"DoNotSchedule"` instructs the scheduler not to schedule the pod when constraints are not satisfied.<br>- `"ScheduleAnyway"` instructs the scheduler to schedule the pod even if constraints are not satisfied. |

## .spec.topologySpreadConstraints[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.topologySpreadConstraints[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.topologySpreadConstraints[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.topologySpreadConstraints[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.topologySpreadConstraints[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be

set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector. This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).

**Type**

`array`

## .spec.topologySpreadConstraints[].matchLabelKeys[]

**Type**

`string`

## .spec.volumes

**Description**

List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

**Type**

`array`

## .spec.volumes[]

**Description**

Volume represents a named volume in a pod that may be accessed by any container in the pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `awsElasticBlockStore` | `object` | Represents a Persistent Disk resource in AWS.<br><br>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling. |
| `azureDisk` | `object` | AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod. |
| `azureFile` | `object` | AzureFile represents an Azure File Service mount on the host and bind mount to the pod. |
| `cephfs` | `object` | Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling. |
| `cinder` | `object` | Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling. |

| Property | Type | Description |
| --- | --- | --- |
| `configMap` | `object` | Adapts a ConfigMap into a volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling. |
| `csi` | `object` | Represents a source location of a volume to mount, managed by an external CSI driver |
| `downwardAPI` | `object` | DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling. |
| `emptyDir` | `object` | Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling. |
| `ephemeral` | `object` | Represents an ephemeral volume that is handled by a normal storage driver. |
| `fc` | `object` | Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling. |
| `flexVolume` | `object` | FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin. |
| `flocker` | `object` | Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling. |
| `gcePersistentDisk` | `object` | Represents a Persistent Disk resource in Google Compute Engine.<br><br>A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling. |
| `gitRepo` | `object` | Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.<br><br>DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container. |
| `glusterfs` | `object` | Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling. |

| Property | Type | Description |
| --- | --- | --- |
| `hostPath` | `object` | Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling. |
| `image` | `object` | ImageVolumeSource represents a image volume resource. |
| `iscsi` | `object` | Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling. |
| `name` | `string` | name of the volume. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `nfs` | `object` | Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling. |
| `persistentVolumeClaim` | `object` | PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system). |
| `photonPersistentDisk` | `object` | Represents a Photon Controller persistent disk resource. |
| `portworxVolume` | `object` | PortworxVolumeSource represents a Portworx volume resource. |
| `projected` | `object` | Represents a projected volume source |
| `quobyte` | `object` | Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling. |
| `rbd` | `object` | Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling. |
| `scaleIO` | `object` | ScaleIOVolumeSource represents a persistent ScaleIO volume |
| `secret` | `object` | Adapts a Secret into a volume.<br><br>The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling. |

| Property | Type | Description |
| --- | --- | --- |
| `storageos` | `object` | Represents a StorageOS persistent volume resource. |
| `vsphereVolume` | `object` | Represents a vSphere volume resource. |

## .spec.volumes[].awsElasticBlockStore

### Description

Represents a Persistent Disk resource in AWS. An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). |
| `readOnly` | `boolean` | readOnly value true will force the readOnly setting in VolumeMounts. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `volumeID` | `string` | volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |

## .spec.volumes[].azureDisk

### Description

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

### Type

`object`

### Required

`diskName` `diskURI`

| Property | Type | Description |
| --- | --- | --- |
| `cachingMode` | `string` | cachingMode is the Host Caching mode: None, Read Only, Read Write.<br>Possible enum values:<br>- `"None"`<br>- `"ReadOnly"`<br>- `"ReadWrite"` |
| `diskName` | `string` | diskName is the Name of the data disk in the blob storage |
| `diskURI` | `string` | diskURI is the URI of data disk in the blob storage |
| `fsType` | `string` | fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `kind` | `string` | kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared<br>Possible enum values:<br>- `"Dedicated"`<br>- `"Managed"`<br>- `"Shared"` |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

## .spec.volumes[].azureFile

### Description

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

### Type

`object`

### Required

`secretName` `shareName`

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretName` | `string` | secretName is the name of secret that contains Azure Storage Account Name and Key |

| Property | Type | Description |
| --- | --- | --- |
| `shareName` | `string` | shareName is the azure share Name |

## .spec.volumes[].cephfs

**Description**

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`monitors`

| Property | Type | Description |
| --- | --- | --- |
| `monitors` | `array` | monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it |
| `path` | `string` | path is Optional: Used as the mounted root, rather than the full Ceph tree, default is / |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it |
| `secretFile` | `string` | secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it |

## .spec.volumes[].cephfs.monitors

**Description**

monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

**Type**

`array`

## .spec.volumes[].cephfs.monitors[]

**Type**

`string`

## .spec.volumes[].cephfs.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.volumes[].cinder

**Description**

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeID` | `string` | volumeID used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |

## .spec.volumes[].cinder.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.volumes[].configMap

### Description

Adapts a ConfigMap into a volume. The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.volumes[].configMap.items

### Description

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

### Type

`array`

## .spec.volumes[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.volumes[].csi

**Description**

Represents a source location of a volume to mount, managed by an external CSI driver

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
|---|---|---|
| `driver` | `string` | driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster. |
| `fsType` | `string` | fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply. |
| `nodePublishSecretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `readOnly` | `boolean` | readOnly specifies a read-only configuration for the volume. Defaults to false (read/write). |

| Property | Type | Description |
| --- | --- | --- |
| `volumeAttributes` | `object` | volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values. |

## .spec.volumes[].csi.nodePublishSecretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.volumes[].csi.volumeAttributes

**Description**

volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**Type**

`object`

## .spec.volumes[].downwardAPI

**Description**

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | Items is a list of downward API volume file |

## .spec.volumes[].downwardAPI.items

**Description**

Items is a list of downward API volume file

**Type**

`array`

## .spec.volumes[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.volumes[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

# .spec.volumes[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

# .spec.volumes[].emptyDir

**Description**

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `medium` | `string` | medium represents what type of storage medium should back this directory. The default is "" which means to use the n |
| `sizeLimit` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAM<br>The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits> <r<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.html)<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a number<br>When a Quantity is parsed from a string, it will remember the type of suffix it had, an<br>Before serializing, Quantity will be put in "canonical form". This means that Exponent/s<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suffix<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point number.<br>Non-canonical values will still parse as long as they are well formed, but will be re-en<br>This format is intended to make it difficult to use these numbers without writing some s |

## .spec.volumes[].ephemeral

**Description**

Represents an ephemeral volume that is handled by a normal storage driver.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `volumeClaimTemplate` | `object` | PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource. |

## .spec.volumes[].ephemeral.volumeClaimTemplate

### Description

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

### Type

`object`

### Required

`spec`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes |

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec

### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `accessModes` | `array` | accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1 ↗ |
| `dataSource` | `object` | TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace. |
| `dataSourceRef` | `object` | TypedObjectReference contains enough information to let you locate the typed referenced object |
| `resources` | `object` | VolumeResourceRequirements describes the storage resource requirements for a volume. |

| Property | Type | Description |
| --- | --- | --- |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `storageClassName` | `string` | storageClassName is the name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1 ↗ |
| `volumeAttributesClassName` | `string` | volumeAttributesClassName may be used to set the VolumeAttributesClass used by this claim. If specified, the CSI driver will create or update the volume with the attributes defined in the corresponding VolumeAttributesClass. This has a different purpose than storageClassName, it can be changed after the claim is created. An empty string value means that no VolumeAttributesClass will be applied to the claim but it's not allowed to reset this field to empty string once it is set. If unspecified and the PersistentVolumeClaim is unbound, the default VolumeAttributesClass will be set by the persistentvolume controller if it exists. If the resource referred to by volumeAttributesClass does not exist, this PersistentVolumeClaim will be set to a Pending state, as reflected by the modifyVolumeStatus field, until such as a resource exists. More info: https://kubernetes.io/docs/concepts/storage/volume-attributes-classes/ ↗ (Beta) Using this field requires the VolumeAttributesClass feature gate to be enabled (off by default). |
| `volumeMode` | `string` | volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.<br><br>Possible enum values:<br>• `"Block"` means the volume will not be formatted with a filesystem and will remain a raw block device.<br>• `"Filesystem"` means the volume will be or is formatted with a filesystem. |
| `volumeName` | `string` | volumeName is the binding reference to the PersistentVolume backing this claim. |

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes

**Description**

accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

**Type**

`array`

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes[]

**Type**

`string`

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

**Description**

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
|---|---|---|
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef

**Description**

TypedObjectReference contains enough information to let you locate the typed referenced object

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
|---|---|---|
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |
| `namespace` | `string` | Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled. |

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

**Description**

VolumeResourceRequirements describes the storage resource requirements for a volume.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.volumes[].fc

**Description**

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `lun` | `integer` | lun is Optional: FC target lun number |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `targetWWNs` | `array` | targetWWNs is Optional: FC target worldwide names (WWNs) |
| `wwids` | `array` | wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously. |

## .spec.volumes[].fc.targetWWNs

**Description**

targetWWNs is Optional: FC target worldwide names (WWNs)

**Type**

`array`

## .spec.volumes[].fc.targetWWNs[]

**Type**

`string`

## .spec.volumes[].fc.wwids

**Description**

wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

**Type**

`array`

## .spec.volumes[].fc.wwids[]

**Type**

`string`

## .spec.volumes[].flexVolume

**Description**

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
| --- | --- | --- |
| `driver` | `string` | driver is the name of the driver to use for this volume. |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script. |
| `options` | `object` | options is Optional: this field holds extra command options if any. |
| `readOnly` | `boolean` | readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |

## .spec.volumes[].flexVolume.options

**Description**

options is Optional: this field holds extra command options if any.

**Type**

`object`

## .spec.volumes[].flexVolume.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.volumes[].flocker

### Description

Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `datasetName` | `string` | datasetName is Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated |
| `datasetUUID` | `string` | datasetUUID is the UUID of the dataset. This is unique identifier of a Flocker dataset |

## .spec.volumes[].gcePersistentDisk

### Description

Represents a Persistent Disk resource in Google Compute Engine. A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

### Type

`object`

### Required

`pdName`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `pdName` | `string` | pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

## .spec.volumes[].gitRepo

**Description**

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

**Type**

`object`

**Required**

`repository`

| Property | Type | Description |
| --- | --- | --- |
| `directory` | `string` | directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name. |
| `repository` | `string` | repository is the URL |
| `revision` | `string` | revision is the commit hash for the specified revision. |

## .spec.volumes[].glusterfs

**Description**

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`endpoints` `path`

| Property | Type | Description |
| --- | --- | --- |
| `endpoints` | `string` | endpoints is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |
| `readOnly` | `boolean` | readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

## .spec.volumes[].hostPath

**Description**

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗ |
| `type` | `string` | type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗<br><br>Possible enum values:<br>• `""` For backwards compatible, leave it empty if unset<br>• `"BlockDevice"` A block device must exist at the given path<br>• `"CharDevice"` A character device must exist at the given path<br>• `"Directory"` A directory must exist at the given path<br>• `"DirectoryOrCreate"` If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet.<br>• `"File"` A file must exist at the given path<br>• `"FileOrCreate"` If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet.<br>• `"Socket"` A UNIX socket must exist at the given path |

## .spec.volumes[].image

**Description**

ImageVolumeSource represents a image volume resource.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `pullPolicy` | `string` | Policy for pulling OCI objects. Possible values are: Always: the kubelet always attempts to pull the reference. Container creation will fail If the pull fails. Never: the kubelet never pulls the reference and only uses a local image or artifact. Container creation will fail if the reference isn't present. IfNotPresent: the kubelet pulls if the reference isn't already present on disk. Container creation will fail if the reference isn't present and the pull fails. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise.<br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `reference` | `string` | Required: Image or artifact reference to be used. Behaves in the same way as pod.spec.containers[*].image. Pull secrets will be assembled in the same way as for the container image by looking up node credentials, SA image pull secrets, and pod spec image pull secrets. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |

## .spec.volumes[].iscsi

### Description

Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`targetPortal` `iqn` `lun`

| Property | Type | Description |
| --- | --- | --- |
| `chapAuthDiscovery` | `boolean` | chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication |
| `chapAuthSession` | `boolean` | chapAuthSession defines whether support iSCSI Session CHAP authentication |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi |
| `initiatorName` | `string` | initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection. |

| Property | Type | Description |
| --- | --- | --- |
| `iqn` | `string` | iqn is the target iSCSI Qualified Name. |
| `iscsiInterface` | `string` | iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp). |
| `lun` | `integer` | lun represents iSCSI Target Lun number. |
| `portals` | `array` | portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `targetPortal` | `string` | targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |

## .spec.volumes[].iscsi.portals

**Description**

portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

**Type**

`array`

## .spec.volumes[].iscsi.portals[]

**Type**

`string`

## .spec.volumes[].iscsi.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.volumes[].nfs

### Description

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`server` `path`

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs ↗ |
| `readOnly` | `boolean` | readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs ↗ |
| `server` | `string` | server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs ↗ |

## .spec.volumes[].persistentVolumeClaim

### Description

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

### Type

`object`

### Required

`claimName`

| Property | Type | Description |
| --- | --- | --- |
| `claimName` | `string` | claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims ↗ |
| `readOnly` | `boolean` | readOnly Will force the ReadOnly setting in VolumeMounts. Default false. |

## .spec.volumes[].photonPersistentDisk

**Description**

Represents a Photon Controller persistent disk resource.

**Type**

`object`

**Required**

`pdID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `pdID` | `string` | pdID is the ID that identifies Photon Controller persistent disk |

## .spec.volumes[].portworxVolume

**Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `volumeID` | `string` | volumeID uniquely identifies a Portworx volume |

## .spec.volumes[].projected

**Description**

Represents a projected volume source

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `sources` | `array` | sources is the list of volume projections. Each entry in this list handles one source. |

## .spec.volumes[].projected.sources

**Description**

sources is the list of volume projections. Each entry in this list handles one source.

**Type**

`array`

## .spec.volumes[].projected.sources[]

**Description**

Projection that may be projected along with other supported volume types. Exactly one of these fields must be set.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `clusterTrustBundle` | `object` | ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem. |
| `configMap` | `object` | Adapts a ConfigMap into a projected volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode. |
| `downwardAPI` | `object` | Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode. |
| `secret` | `object` | Adapts a secret into a projected volume.<br><br>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode. |
| `serviceAccountToken` | `object` | ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API |

| Property | Type | Description |
| --- | --- | --- |
| | | Server or otherwise). |

## .spec.volumes[].projected.sources[].clusterTrustBundle

**Description**

ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem.

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `name` | `string` | Select a single ClusterTrustBundle by object name. Mutually-exclusive with signerName and labelSelector. |
| `optional` | `boolean` | If true, don't block pod startup if the referenced ClusterTrustBundle(s) aren't available. If using name, then the named ClusterTrustBundle is allowed not to exist. If using signerName, then the combination of signerName and labelSelector is allowed to match zero ClusterTrustBundles. |
| `path` | `string` | Relative path from the volume root to write the bundle. |
| `signerName` | `string` | Select all ClusterTrustBundles that match this signer name. Mutually-exclusive with name. The contents of all selected ClusterTrustBundles will be unified and deduplicated. |

## .spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only |

| Property | Type | Description |
| --- | --- | --- |
| | | "value". The requirements are ANDed. |

## .spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.volumes[].projected.sources[].configMap

**Description**

Adapts a ConfigMap into a projected volume. The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.volumes[].projected.sources[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.volumes[].projected.sources[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.volumes[].projected.sources[].downwardAPI

**Description**

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `items` | `array` | Items is a list of DownwardAPIVolume file |

## .spec.volumes[].projected.sources[].downwardAPI.items

**Description**

Items is a list of DownwardAPIVolume file

**Type**

`array`

## .spec.volumes[].projected.sources[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |

| Property | Type | Description |
| --- | --- | --- |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is:<br><br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.volumes[].projected.sources[].secret

**Description**

Adapts a secret into a projected volume. The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional field specify whether the Secret or its key must be defined |

## .spec.volumes[].projected.sources[].secret.items

### Description

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

### Type

`array`

## .spec.volumes[].projected.sources[].secret.items[]

### Description

Maps a string key to a path within a volume.

### Type

`object`

### Required

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.volumes[].projected.sources[].serviceAccountToken

**Description**

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `audience` | `string` | audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver. |
| `expirationSeconds` | `integer` | expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours.Defaults to 1 hour and must be at least 10 minutes. |
| `path` | `string` | path is the path relative to the mount point of the file to project the token into. |

## .spec.volumes[].quobyte

**Description**

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`registry` `volume`

| Property | Type | Description |
|---|---|---|
| `group` | `string` | group to map volume access to Default is no group |
| `readOnly` | `boolean` | readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false. |
| `registry` | `string` | registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes |

| Property | Type | Description |
| --- | --- | --- |
| `tenant` | `string` | tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin |
| `user` | `string` | user to map volume access to Defaults to serivceaccount user |
| `volume` | `string` | volume is a string that references an already created Quobyte volume by name. |

## .spec.volumes[].rbd

### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`monitors` `image`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd ↗ |
| `image` | `string` | image is the rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `keyring` | `string` | keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `monitors` | `array` | monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `pool` | `string` | pool is the rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |

| Property | Type | Description |
| --- | --- | --- |
| `user` | `string` | user is the rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |

## .spec.volumes[].rbd.monitors

**Description**

monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

**Type**

`array`

## .spec.volumes[].rbd.monitors[]

**Type**

`string`

## .spec.volumes[].rbd.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.volumes[].scaleIO

**Description**

ScaleIOVolumeSource represents a persistent ScaleIO volume

**Type**

`object`

**Required**

`gateway` `system` `secretRef`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs". |

| Property | Type | Description |
| --- | --- | --- |
| `gateway` | `string` | gateway is the host address of the ScaleIO API Gateway. |
| `protectionDomain` | `string` | protectionDomain is the name of the ScaleIO Protection Domain for the configured storage. |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `sslEnabled` | `boolean` | sslEnabled Flag enable/disable SSL communication with Gateway, default false |
| `storageMode` | `string` | storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned. |
| `storagePool` | `string` | storagePool is the ScaleIO Storage Pool associated with the protection domain. |
| `system` | `string` | system is the name of the storage system as configured in ScaleIO. |
| `volumeName` | `string` | volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source. |

## .spec.volumes[].scaleIO.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.volumes[].secret

**Description**

Adapts a Secret into a volume. The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `defaultMode` | `integer` | defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `optional` | `boolean` | optional field specify whether the Secret or its keys must be defined |
| `secretName` | `string` | secretName is the name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret ↗ |

## .spec.volumes[].secret.items

**Description**

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.volumes[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.volumes[].storageos

**Description**

Represents a StorageOS persistent volume resource.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeName` | `string` | volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace. |
| `volumeNamespace` | `string` | volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created. |

## .spec.volumes[].storageos.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.volumes[].vsphereVolume

### Description

Represents a vSphere volume resource.

### Type

`object`

### Required

`volumePath`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `storagePolicyID` | `string` | storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName. |
| `storagePolicyName` | `string` | storagePolicyName is the storage Policy Based Management (SPBM) profile name. |
| `volumePath` | `string` | volumePath is the path that identifies vSphere volume vmdk |

## .status

### Description

PodStatus represents information about the status of a pod. Status may trail the actual state of a system, especially if the node that hosts the pod cannot contact the control plane.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `conditions` | `array` | Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `containerStatuses` | `array` | Statuses of containers in this pod. Each container in the pod should have at most one status in this list, and all statuses should be for containers in the pod. However this is not enforced. If a status for a non-existent container is present in the list, or the list has duplicate names, the behavior of various Kubernetes components is not defined and those statuses might be ignored. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status ↗ |
| `ephemeralContainerStatuses` | `array` | Statuses for any ephemeral containers that have run in this pod. Each ephemeral container in the pod should have at most one status in this list, and all statuses should be for containers in the pod. However this is not enforced. If a status for a non-existent container is present in the list, or the list has duplicate names, the behavior of various Kubernetes components is not defined and those statuses might be ignored. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status ↗ |
| `hostIP` | `string` | hostIP holds the IP address of the host to which the pod is assigned. Empty if the pod has not started yet. A pod can be assigned to a node that has a problem in kubelet which in turns mean that HostIP will not be updated even if there is a node is assigned to pod |
| `hostIPs` | `array` | hostIPs holds the IP addresses allocated to the host. If this field is specified, the first entry must match the hostIP field. This list is empty if the pod has not started yet. A pod can be assigned to a node that has a problem in kubelet which in turns means that HostIPs will not be updated even if there is a node is assigned to this pod. |
| `initContainerStatuses` | `array` | Statuses of init containers in this pod. The most recent successful non-restartable init container will have ready = true, the most recently started container will have startTime set. Each init container in the pod should have at most one status in this list, and all statuses should be for containers in the pod. However this is not enforced. If a status for a non-existent container is present in the list, or the list has duplicate names, the behavior of various Kubernetes components is not defined and those statuses might be ignored. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#pod-and-container-status ↗ |
| `message` | `string` | A human readable message indicating details about why the pod is in this condition. |
| `nominatedNodeName` | `string` | nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled. |
| `phase` | `string` | The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: |

| Property | Type | Description |
| --- | --- | --- |
| | | Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod.<br><br>More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase ↗<br><br>Possible enum values:<br><br>• `"Failed"` means that all containers in the pod have terminated, and at least one container has terminated in a failure (exited with a non-zero exit code or was stopped by the system).<br>• `"Pending"` means the pod has been accepted by the system, but one or more of the containers has not been started. This includes time before being bound to a node, as well as time spent pulling images onto the host.<br>• `"Running"` means the pod has been bound to a node and all of the containers have been started. At least one container is still running or is in the process of being restarted.<br>• `"Succeeded"` means that all containers in the pod have voluntarily terminated with a container exit code of 0, and the system is not going to restart any of these containers.<br>• `"Unknown"` means that for some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. Deprecated: It isn't being set since 2015 (74da3b14b0c0f658b3bb8d2def5094686d0e9095) |
| `podIP` | `string` | podIP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated. |
| `podIPs` | `array` | podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet. |
| `qosClass` | `string` | The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-qos/#quality-of-service-classes ↗<br><br>Possible enum values:<br><br>• `"BestEffort"` is the BestEffort qos class.<br>• `"Burstable"` is the Burstable qos class.<br>• `"Guaranteed"` is the Guaranteed qos class. |
| `reason` | `string` | A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted' |

| Property | Type | Description |
| --- | --- | --- |
| `resize` | `string` | Status of resources resize desired for pod's containers. It is empty if no resources resize is pending. Any changes to container resources will automatically set this to "Proposed" |
| `resourceClaimStatuses` | `array` | Status of resource claims. |
| `startTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.conditions

### Description

Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

### Type

`array`

## .status.conditions[]

### Description

PodCondition contains details for the current condition of this pod.

### Type

`object`

### Required

`type` `status`

| Property | Type | Description |
| --- | --- | --- |
| `lastProbeTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `lastTransitionTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | Human-readable message indicating details about last transition. |
| `reason` | `string` | Unique, one-word, CamelCase reason for the condition's last transition. |
| `status` | `string` | Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `type` | `string` | Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions |

## .status.containerStatuses

### Description

Statuses of containers in this pod. Each container in the pod should have at most one status in this list, and all statuses should be for containers in the pod. However this is not enforced. If a status for a non-existent container is present in the list, or the list has duplicate names, the behavior of various Kubernetes components is not defined and those statuses might be ignored. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status

### Type

`array`

## .status.containerStatuses[]

### Description

ContainerStatus contains details for the current status of this container.

### Type

`object`

### Required

`name` `ready` `restartCount` `image` `imageID`

| Property | Type | Description |
| --- | --- | --- |
| `allocatedResources` | `object` | AllocatedResources represents the compute resources allocated for this container by the node. Kubelet sets this value to Container.Resources.Requests upon successful pod admission and after successfully admitting desired pod resize. |
| `allocatedResourcesStatus` | `array` | AllocatedResourcesStatus represents the status of various resources allocated for this Pod. |
| `containerID` | `string` | ContainerID is the ID of the container in the format '://<container_id>'. Where type is a container runtime identifier, returned from Version call of CRI API (for example "containerd"). |
| `image` | `string` | Image is the name of container image that the container is running. The container image may not match the image used in the PodSpec, as it may have been resolved by the runtime. More info: https://kubernetes.io/docs/concepts/containers/images. |
| `imageID` | `string` | ImageID is the image ID of the container's image. The image ID may not match the image ID of the image used in the PodSpec, as it may have been resolved by the runtime. |

| Property | Type | Description |
| --- | --- | --- |
| `lastState` | `object` | ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting. |
| `name` | `string` | Name is a DNS_LABEL representing the unique name of the container. Each container in a pod must have a unique name across all container types. Cannot be updated. |
| `ready` | `boolean` | Ready specifies whether the container is currently passing its readiness check. The value will change as readiness probes keep executing. If no readiness probes are specified, this field defaults to true once the container is fully started (see Started field).<br><br>The value is typically used to determine whether a container is ready to accept traffic. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartCount` | `integer` | RestartCount holds the number of times the container has been restarted. Kubelet makes an effort to always increment the value, but there are cases when the state may be lost due to node restarts and then the value may be reset to 0. The value is never negative. |
| `started` | `boolean` | Started indicates whether the container has finished its postStart lifecycle hook and passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. In both cases, startup probes will run again. Is always true when no startupProbe is defined and container is running and has passed the postStart lifecycle hook. The null value must be treated the same as false. |
| `state` | `object` | ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting. |
| `user` | `object` | ContainerUser represents user identity information |
| `volumeMounts` | `array` | Status of volume mounts. |

## .status.containerStatuses[].allocatedResources

**Description**

AllocatedResources represents the compute resources allocated for this container by the node. Kubelet sets this value to Container.Resources.Requests upon successful pod admission and after successfully admitting desired pod resize.

**Type**

`object`

## .status.containerStatuses[].allocatedResourcesStatus

**Description**

AllocatedResourcesStatus represents the status of various resources allocated for this Pod.

**Type**

`array`

## .status.containerStatuses[].allocatedResourcesStatus[]

**Description**

ResourceStatus represents the status of a single resource allocated to a Pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the resource. Must be unique within the pod and in case of non-DRA resource, match one of the resources from the pod spec. For DRA resources, the value must be "claim:<claim_name>/". When this status is reported about a container, the "claim_name" and "request" must match one of the claims of this container. |
| `resources` | `array` | List of unique resources health. Each element in the list contains an unique resource ID and its health. At a minimum, for the lifetime of a Pod, resource ID must uniquely identify the resource allocated to the Pod on the Node. If other Pod on the same Node reports the status with the same resource ID, it must be the same resource they share. See ResourceID type definition for a specific format it has in various use cases. |

## .status.containerStatuses[].allocatedResourcesStatus[].resources

**Description**

List of unique resources health. Each element in the list contains an unique resource ID and its health. At a minimum, for the lifetime of a Pod, resource ID must uniquely identify the resource allocated to the Pod on the Node. If other Pod on the same Node reports the status with the same resource ID, it must be the same resource they share. See ResourceID type definition for a specific format it has in various use cases.

**Type**

`array`

## .status.containerStatuses[].allocatedResourcesStatus[].resources[]

**Description**

ResourceHealth represents the health of a resource. It has the latest device health information. This is a part of KEP https://kep.k8s.io/4680.

**Type**

`object`

**Required**

`resourceID`

| Property | Type | Description |
| --- | --- | --- |
| `health` | `string` | Health of the resource. can be one of:<br>• Healthy: operates as normal<br>• Unhealthy: reported unhealthy. We consider this a temporary health issue since we do not have a mechanism today to distinguish temporary and permanent issues.<br>• Unknown: The status cannot be determined. For example, Device Plugin got unregistered and hasn't been re-registered since.<br>In future we may want to introduce the PermanentlyUnhealthy Status. |
| `resourceID` | `string` | ResourceID is the unique identifier of the resource. See the ResourceID type for more information. |

## .status.containerStatuses[].lastState

**Description**

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `running` | `object` | ContainerStateRunning is a running state of a container. |
| `terminated` | `object` | ContainerStateTerminated is a terminated state of a container. |
| `waiting` | `object` | ContainerStateWaiting is a waiting state of a container. |

## .status.containerStatuses[].lastState.running

**Description**

ContainerStateRunning is a running state of a container.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.containerStatuses[].lastState.terminated

**Description**

ContainerStateTerminated is a terminated state of a container.

**Type**

`object`

**Required**

`exitCode`

| Property | Type | Description |
|---|---|---|
| `containerID` | `string` | Container's ID in the format '://<container_id>' |
| `exitCode` | `integer` | Exit status from the last termination of the container |
| `finishedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | Message regarding the last termination of the container |
| `reason` | `string` | (brief) reason from the last termination of the container |
| `signal` | `integer` | Signal from the last termination of the container |
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.containerStatuses[].lastState.waiting

**Description**

ContainerStateWaiting is a waiting state of a container.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `message` | `string` | Message regarding why the container is not yet running. |
| `reason` | `string` | (brief) reason the container is not yet running. |

## .status.containerStatuses[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .status.containerStatuses[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .status.containerStatuses[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .status.containerStatuses[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .status.containerStatuses[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .status.containerStatuses[].state

**Description**

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `running` | `object` | ContainerStateRunning is a running state of a container. |
| `terminated` | `object` | ContainerStateTerminated is a terminated state of a container. |
| `waiting` | `object` | ContainerStateWaiting is a waiting state of a container. |

## .status.containerStatuses[].state.running

**Description**

ContainerStateRunning is a running state of a container.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.containerStatuses[].state.terminated

**Description**

ContainerStateTerminated is a terminated state of a container.

**Type**

`object`

**Required**

`exitCode`

| Property | Type | Description |
|---|---|---|
| `containerID` | `string` | Container's ID in the format '://<container_id>' |
| `exitCode` | `integer` | Exit status from the last termination of the container |
| `finishedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | Message regarding the last termination of the container |
| `reason` | `string` | (brief) reason from the last termination of the container |
| `signal` | `integer` | Signal from the last termination of the container |
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.containerStatuses[].state.waiting

**Description**

ContainerStateWaiting is a waiting state of a container.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `message` | `string` | Message regarding why the container is not yet running. |
| `reason` | `string` | (brief) reason the container is not yet running. |

## .status.containerStatuses[].user

**Description**

ContainerUser represents user identity information

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `linux` | `object` | LinuxContainerUser represents user identity information in Linux containers |

## .status.containerStatuses[].user.linux

**Description**

LinuxContainerUser represents user identity information in Linux containers

**Type**

`object`

**Required**

`uid` `gid`

| Property | Type | Description |
| --- | --- | --- |
| `gid` | `integer` | GID is the primary gid initially attached to the first process in the container |
| `supplementalGroups` | `array` | SupplementalGroups are the supplemental groups initially attached to the first process in the container |
| `uid` | `integer` | UID is the primary uid initially attached to the first process in the container |

## .status.containerStatuses[].user.linux.supplementalGroups

**Description**

SupplementalGroups are the supplemental groups initially attached to the first process in the container

**Type**

`array`

## .status.containerStatuses[].user.linux.supplementalGroups[]

**Type**

`integer`

## .status.containerStatuses[].volumeMounts

**Description**

Status of volume mounts.

**Type**

`array`

## .status.containerStatuses[].volumeMounts[]

**Description**

VolumeMountStatus shows status of volume mounts.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | MountPath corresponds to the original VolumeMount. |
| `name` | `string` | Name corresponds to the name of the original VolumeMount. |
| `readOnly` | `boolean` | ReadOnly corresponds to the original VolumeMount. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly must be set to Disabled, Enabled, or unspecified (for non-readonly mounts). An IfPossible value in the original VolumeMount must be translated to Disabled or Enabled, depending on the mount result. |

## .status.ephemeralContainerStatuses

**Description**

Statuses for any ephemeral containers that have run in this pod. Each ephemeral container in the pod should have at most one status in this list, and all statuses should be for containers in the pod. However this is not enforced. If a status for a non-existent container is present in the list, or the list has duplicate names, the behavior of various Kubernetes components is not defined and those statuses might be ignored. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status

**Type**

`array`

## .status.ephemeralContainerStatuses[]

**Description**

ContainerStatus contains details for the current status of this container.

**Type**

`object`

**Required**

`name` `ready` `restartCount` `image` `imageID`

| Property | Type | Description |
| --- | --- | --- |
| `allocatedResources` | `object` | AllocatedResources represents the compute resources allocated for this container by the node. Kubelet sets this value to Container.Resources.Requests upon successful pod admission and after successfully admitting desired pod resize. |
| `allocatedResourcesStatus` | `array` | AllocatedResourcesStatus represents the status of various resources allocated for this Pod. |
| `containerID` | `string` | ContainerID is the ID of the container in the format '://<container_id>'. Where type is a container runtime identifier, returned from Version call of CRI API (for example "containerd"). |
| `image` | `string` | Image is the name of container image that the container is running. The container image may not match the image used in the PodSpec, as it may have been resolved by the runtime. More info: https://kubernetes.io/docs/concepts/containers/images ↗. |
| `imageID` | `string` | ImageID is the image ID of the container's image. The image ID may not match the image ID of the image used in the PodSpec, as it may have been resolved by the runtime. |
| `lastState` | `object` | ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting. |
| `name` | `string` | Name is a DNS_LABEL representing the unique name of the container. Each container in a pod must have a unique name across all container types. Cannot be updated. |
| `ready` | `boolean` | Ready specifies whether the container is currently passing its readiness check. The value will change as readiness probes keep executing. If no readiness probes are specified, this field defaults to true once the container is fully started (see Started field).<br><br>The value is typically used to determine whether a container is ready to accept traffic. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartCount` | `integer` | RestartCount holds the number of times the container has been restarted. Kubelet makes an effort to always increment the value, but there are cases when the state may be lost due to node restarts and then the value may be reset to 0. The value is never negative. |

| Property | Type | Description |
| --- | --- | --- |
| `started` | `boolean` | Started indicates whether the container has finished its postStart lifecycle hook and passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. In both cases, startup probes will run again. Is always true when no startupProbe is defined and container is running and has passed the postStart lifecycle hook. The null value must be treated the same as false. |
| `state` | `object` | ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting. |
| `user` | `object` | ContainerUser represents user identity information |
| `volumeMounts` | `array` | Status of volume mounts. |

## .status.ephemeralContainerStatuses[].allocatedResources

**Description**

AllocatedResources represents the compute resources allocated for this container by the node. Kubelet sets this value to Container.Resources.Requests upon successful pod admission and after successfully admitting desired pod resize.

**Type**

`object`

## .status.ephemeralContainerStatuses[].allocatedResourcesStatus

**Description**

AllocatedResourcesStatus represents the status of various resources allocated for this Pod.

**Type**

`array`

## .status.ephemeralContainerStatuses[].allocatedResourcesStatus[]

**Description**

ResourceStatus represents the status of a single resource allocated to a Pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the resource. Must be unique within the pod and in case of non-DRA resource, match one of the resources from the pod spec. For DRA resources, the value must be "claim:<claim_name>/". When this status is reported about a container, the "claim_name" and "request" must match one of the claims of this container. |
| `resources` | `array` | List of unique resources health. Each element in the list contains an unique resource ID and its health. At a minimum, for the lifetime of a Pod, resource ID must uniquely identify the resource allocated to the Pod on the Node. If other Pod on the same Node reports the status with the same resource ID, it must be the same resource they share. See ResourceID type definition for a specific format it has in various use cases. |

## .status.ephemeralContainerStatuses[].allocatedResourcesStatus[].resources

**Description**

List of unique resources health. Each element in the list contains an unique resource ID and its health. At a minimum, for the lifetime of a Pod, resource ID must uniquely identify the resource allocated to the Pod on the Node. If other Pod on the same Node reports the status with the same resource ID, it must be the same resource they share. See ResourceID type definition for a specific format it has in various use cases.

**Type**

`array`

## .status.ephemeralContainerStatuses[].allocatedResourcesStatus[].resources[]

**Description**

ResourceHealth represents the health of a resource. It has the latest device health information. This is a part of KEP https://kep.k8s.io/4680.

**Type**

`object`

**Required**

`resourceID`

| Property | Type | Description |
| --- | --- | --- |
| `health` | `string` | Health of the resource. can be one of:<br>• Healthy: operates as normal<br>• Unhealthy: reported unhealthy. We consider this a temporary health issue since we do not have a mechanism today to distinguish temporary and permanent issues.<br>• Unknown: The status cannot be determined. For example, Device Plugin got unregistered and hasn't been re-registered since.<br>In future we may want to introduce the PermanentlyUnhealthy Status. |
| `resourceID` | `string` | ResourceID is the unique identifier of the resource. See the ResourceID type for more information. |

## .status.ephemeralContainerStatuses[].lastState

**Description**

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `running` | `object` | ContainerStateRunning is a running state of a container. |
| `terminated` | `object` | ContainerStateTerminated is a terminated state of a container. |
| `waiting` | `object` | ContainerStateWaiting is a waiting state of a container. |

## .status.ephemeralContainerStatuses[].lastState.running

**Description**

ContainerStateRunning is a running state of a container.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.ephemeralContainerStatuses[].lastState.terminated

**Description**

ContainerStateTerminated is a terminated state of a container.

**Type**

`object`

**Required**

`exitCode`

| Property | Type | Description |
|---|---|---|
| `containerID` | `string` | Container's ID in the format '://<container_id>' |
| `exitCode` | `integer` | Exit status from the last termination of the container |

| Property | Type | Description |
| --- | --- | --- |
| `finishedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | Message regarding the last termination of the container |
| `reason` | `string` | (brief) reason from the last termination of the container |
| `signal` | `integer` | Signal from the last termination of the container |
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.ephemeralContainerStatuses[].lastState.waiting

**Description**

ContainerStateWaiting is a waiting state of a container.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `message` | `string` | Message regarding why the container is not yet running. |
| `reason` | `string` | (brief) reason the container is not yet running. |

## .status.ephemeralContainerStatuses[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .status.ephemeralContainerStatuses[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .status.ephemeralContainerStatuses[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .status.ephemeralContainerStatuses[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .status.ephemeralContainerStatuses[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .status.ephemeralContainerStatuses[].state

**Description**

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `running` | `object` | ContainerStateRunning is a running state of a container. |
| `terminated` | `object` | ContainerStateTerminated is a terminated state of a container. |
| `waiting` | `object` | ContainerStateWaiting is a waiting state of a container. |

## .status.ephemeralContainerStatuses[].state.running

**Description**

ContainerStateRunning is a running state of a container.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.ephemeralContainerStatuses[].state.terminated

**Description**

ContainerStateTerminated is a terminated state of a container.

**Type**

`object`

**Required**

`exitCode`

| Property | Type | Description |
| --- | --- | --- |
| `containerID` | `string` | Container's ID in the format '://<container_id>' |
| `exitCode` | `integer` | Exit status from the last termination of the container |
| `finishedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | Message regarding the last termination of the container |
| `reason` | `string` | (brief) reason from the last termination of the container |
| `signal` | `integer` | Signal from the last termination of the container |
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.ephemeralContainerStatuses[].state.waiting

**Description**

ContainerStateWaiting is a waiting state of a container.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `message` | `string` | Message regarding why the container is not yet running. |
| `reason` | `string` | (brief) reason the container is not yet running. |

## .status.ephemeralContainerStatuses[].user

**Description**

ContainerUser represents user identity information

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `linux` | `object` | LinuxContainerUser represents user identity information in Linux containers |

## .status.ephemeralContainerStatuses[].user.linux

**Description**

LinuxContainerUser represents user identity information in Linux containers

**Type**

`object`

**Required**

`uid` `gid`

| Property | Type | Description |
| --- | --- | --- |
| `gid` | `integer` | GID is the primary gid initially attached to the first process in the container |
| `supplementalGroups` | `array` | SupplementalGroups are the supplemental groups initially attached to the first process in the container |
| `uid` | `integer` | UID is the primary uid initially attached to the first process in the container |

## .status.ephemeralContainerStatuses[].user.linux.supplementalGroups

**Description**

SupplementalGroups are the supplemental groups initially attached to the first process in the container

**Type**

`array`

## .status.ephemeralContainerStatuses[].user.linux.supplementalGroups[]

**Type**

`integer`

## .status.ephemeralContainerStatuses[].volumeMounts

**Description**

Status of volume mounts.

**Type**

`array`

## .status.ephemeralContainerStatuses[].volumeMounts[]

**Description**

VolumeMountStatus shows status of volume mounts.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
|---|---|---|
| `mountPath` | `string` | MountPath corresponds to the original VolumeMount. |
| `name` | `string` | Name corresponds to the name of the original VolumeMount. |
| `readOnly` | `boolean` | ReadOnly corresponds to the original VolumeMount. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly must be set to Disabled, Enabled, or unspecified (for non-readonly mounts). An IfPossible value in the original VolumeMount must be translated to Disabled or Enabled, depending on the mount result. |

## .status.hostIPs

**Description**

hostIPs holds the IP addresses allocated to the host. If this field is specified, the first entry must match the hostIP field. This list is empty if the pod has not started yet. A pod can be assigned to a node that has a problem in kubelet which in turns means that HostIPs will not be updated even if there is a node is assigned to this pod.

**Type**

`array`

## .status.hostIPs[]

**Description**

HostIP represents a single IP address allocated to the host.

**Type**

`object`

**Required**

`ip`

| Property | Type | Description |
|---|---|---|
| `ip` | `string` | IP is the IP address assigned to the host |

## .status.initContainerStatuses

**Description**

Statuses of init containers in this pod. The most recent successful non-restartable init container will have ready = true, the most recently started container will have startTime set. Each init container in the pod should have at most one status in this list, and all statuses should be for containers in the pod. However this is not enforced. If a status for a non-existent container is present in the list, or the list has duplicate names, the behavior of various Kubernetes components is not defined and those statuses might be ignored. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#pod-and-container-status

**Type**

`array`

## .status.initContainerStatuses[]

**Description**

ContainerStatus contains details for the current status of this container.

**Type**

`object`

**Required**

`name` `ready` `restartCount` `image` `imageID`

| Property | Type | Description |
|---|---|---|
| `allocatedResources` | `object` | AllocatedResources represents the compute resources allocated for this container by the node. Kubelet sets this value to Container.Resources.Requests upon successful pod admission and after successfully admitting desired pod resize. |
| `allocatedResourcesStatus` | `array` | AllocatedResourcesStatus represents the status of various resources allocated for this Pod. |
| `containerID` | `string` | ContainerID is the ID of the container in the format '://<container_id>'. Where type is a container runtime identifier, returned from Version call of CRI API (for example "containerd"). |
| `image` | `string` | Image is the name of container image that the container is running. The container image may not match the image used in the PodSpec, as it may have been resolved by the runtime. More info: https://kubernetes.io/docs/concepts/containers/images ↗. |
| `imageID` | `string` | ImageID is the image ID of the container's image. The image ID may not match the image ID of the image used in the PodSpec, as it may have been resolved by the runtime. |
| `lastState` | `object` | ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting. |
| `name` | `string` | Name is a DNS_LABEL representing the unique name of the container. Each container in a pod must have a unique name across all container types. Cannot be updated. |
| `ready` | `boolean` | Ready specifies whether the container is currently passing its readiness check. The value will change as readiness probes keep executing. If no readiness probes are specified, this field defaults to true |

| Property | Type | Description |
| --- | --- | --- |
| | | once the container is fully started (see Started field).<br><br>The value is typically used to determine whether a container is ready to accept traffic. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartCount` | `integer` | RestartCount holds the number of times the container has been restarted. Kubelet makes an effort to always increment the value, but there are cases when the state may be lost due to node restarts and then the value may be reset to 0. The value is never negative. |
| `started` | `boolean` | Started indicates whether the container has finished its postStart lifecycle hook and passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. In both cases, startup probes will run again. Is always true when no startupProbe is defined and container is running and has passed the postStart lifecycle hook. The null value must be treated the same as false. |
| `state` | `object` | ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting. |
| `user` | `object` | ContainerUser represents user identity information |
| `volumeMounts` | `array` | Status of volume mounts. |

## .status.initContainerStatuses[].allocatedResources

**Description**

AllocatedResources represents the compute resources allocated for this container by the node. Kubelet sets this value to Container.Resources.Requests upon successful pod admission and after successfully admitting desired pod resize.

**Type**

`object`

## .status.initContainerStatuses[].allocatedResourcesStatus

**Description**

AllocatedResourcesStatus represents the status of various resources allocated for this Pod.

**Type**

`array`

## .status.initContainerStatuses[].allocatedResourcesStatus[]

**Description**

ResourceStatus represents the status of a single resource allocated to a Pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the resource. Must be unique within the pod and in case of non-DRA resource, match one of the resources from the pod spec. For DRA resources, the value must be "claim:<claim_name>/". When this status is reported about a container, the "claim_name" and "request" must match one of the claims of this container. |
| `resources` | `array` | List of unique resources health. Each element in the list contains an unique resource ID and its health. At a minimum, for the lifetime of a Pod, resource ID must uniquely identify the resource allocated to the Pod on the Node. If other Pod on the same Node reports the status with the same resource ID, it must be the same resource they share. See ResourceID type definition for a specific format it has in various use cases. |

## .status.initContainerStatuses[].allocatedResourcesStatus[].resources

**Description**

List of unique resources health. Each element in the list contains an unique resource ID and its health. At a minimum, for the lifetime of a Pod, resource ID must uniquely identify the resource allocated to the Pod on the Node. If other Pod on the same Node reports the status with the same resource ID, it must be the same resource they share. See ResourceID type definition for a specific format it has in various use cases.

**Type**

`array`

## .status.initContainerStatuses[].allocatedResourcesStatus[].resources[]

**Description**

ResourceHealth represents the health of a resource. It has the latest device health information. This is a part of KEP https://kep.k8s.io/4680.

**Type**

`object`

**Required**

`resourceID`

| Property | Type | Description |
|---|---|---|
| `health` | `string` | Health of the resource. can be one of:<br>• Healthy: operates as normal<br>• Unhealthy: reported unhealthy. We consider this a temporary health issue since we do not have a mechanism today to distinguish temporary and permanent issues.<br>• Unknown: The status cannot be determined. For example, Device Plugin got unregistered and hasn't been re-registered since.<br>In future we may want to introduce the PermanentlyUnhealthy Status. |

| Property | Type | Description |
| --- | --- | --- |
| `resourceID` | `string` | ResourceID is the unique identifier of the resource. See the ResourceID type for more information. |

## .status.initContainerStatuses[].lastState

**Description**

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `running` | `object` | ContainerStateRunning is a running state of a container. |
| `terminated` | `object` | ContainerStateTerminated is a terminated state of a container. |
| `waiting` | `object` | ContainerStateWaiting is a waiting state of a container. |

## .status.initContainerStatuses[].lastState.running

**Description**

ContainerStateRunning is a running state of a container.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.initContainerStatuses[].lastState.terminated

**Description**

ContainerStateTerminated is a terminated state of a container.

**Type**

`object`

**Required**

`exitCode`

| Property | Type | Description |
| --- | --- | --- |
| `containerID` | `string` | Container's ID in the format '://<container_id>' |
| `exitCode` | `integer` | Exit status from the last termination of the container |
| `finishedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | Message regarding the last termination of the container |
| `reason` | `string` | (brief) reason from the last termination of the container |
| `signal` | `integer` | Signal from the last termination of the container |
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.initContainerStatuses[].lastState.waiting

**Description**

ContainerStateWaiting is a waiting state of a container.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `message` | `string` | Message regarding why the container is not yet running. |
| `reason` | `string` | (brief) reason the container is not yet running. |

## .status.initContainerStatuses[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br><br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .status.initContainerStatuses[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .status.initContainerStatuses[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .status.initContainerStatuses[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .status.initContainerStatuses[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .status.initContainerStatuses[].state

**Description**

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `running` | `object` | ContainerStateRunning is a running state of a container. |
| `terminated` | `object` | ContainerStateTerminated is a terminated state of a container. |
| `waiting` | `object` | ContainerStateWaiting is a waiting state of a container. |

## .status.initContainerStatuses[].state.running

**Description**

ContainerStateRunning is a running state of a container.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.initContainerStatuses[].state.terminated

**Description**

ContainerStateTerminated is a terminated state of a container.

**Type**

`object`

**Required**

`exitCode`

| Property | Type | Description |
| --- | --- | --- |
| `containerID` | `string` | Container's ID in the format '://<container_id>' |
| `exitCode` | `integer` | Exit status from the last termination of the container |
| `finishedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | Message regarding the last termination of the container |
| `reason` | `string` | (brief) reason from the last termination of the container |
| `signal` | `integer` | Signal from the last termination of the container |
| `startedAt` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |

## .status.initContainerStatuses[].state.waiting

**Description**

ContainerStateWaiting is a waiting state of a container.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `message` | `string` | Message regarding why the container is not yet running. |
| `reason` | `string` | (brief) reason the container is not yet running. |

## .status.initContainerStatuses[].user

**Description**

ContainerUser represents user identity information

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `linux` | `object` | LinuxContainerUser represents user identity information in Linux containers |

## .status.initContainerStatuses[].user.linux

**Description**

LinuxContainerUser represents user identity information in Linux containers

**Type**

`object`

**Required**

`uid` `gid`

| Property | Type | Description |
| --- | --- | --- |
| `gid` | `integer` | GID is the primary gid initially attached to the first process in the container |
| `supplementalGroups` | `array` | SupplementalGroups are the supplemental groups initially attached to the first process in the container |
| `uid` | `integer` | UID is the primary uid initially attached to the first process in the container |

## .status.initContainerStatuses[].user.linux.supplementalGroups

**Description**

SupplementalGroups are the supplemental groups initially attached to the first process in the container

**Type**

`array`

## .status.initContainerStatuses[].user.linux.supplementalGroups[]

**Type**

`integer`

## .status.initContainerStatuses[].volumeMounts

**Description**

Status of volume mounts.

**Type**

`array`

## .status.initContainerStatuses[].volumeMounts[]

**Description**

VolumeMountStatus shows status of volume mounts.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
|---|---|---|
| `mountPath` | `string` | MountPath corresponds to the original VolumeMount. |
| `name` | `string` | Name corresponds to the name of the original VolumeMount. |
| `readOnly` | `boolean` | ReadOnly corresponds to the original VolumeMount. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly must be set to Disabled, Enabled, or unspecified (for non-readonly mounts). An IfPossible value in the original VolumeMount must be translated to Disabled or Enabled, depending on the mount result. |

## .status.podIPs

**Description**

podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

**Type**

`array`

## .status.podIPs[]

**Description**

PodIP represents a single IP address allocated to the pod.

**Type**

`object`

**Required**

`ip`

| Property | Type | Description |
|---|---|---|
| `ip` | `string` | IP is the IP address assigned to the pod |

## .status.resourceClaimStatuses

**Description**

Status of resource claims.

**Type**

`array`

## .status.resourceClaimStatuses[]

**Description**

PodResourceClaimStatus is stored in the PodStatus for each PodResourceClaim which references a ResourceClaimTemplate. It stores the generated name for the corresponding ResourceClaim.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name uniquely identifies this resource claim inside the pod. This must match the name of an entry in pod.spec.resourceClaims, which implies that the string must be a DNS_LABEL. |
| `resourceClaimName` | `string` | ResourceClaimName is the name of the ResourceClaim that was generated for the Pod in the namespace of the Pod. If this is unset, then generating a ResourceClaim was not necessary. The pod.spec.resourceClaims entry can be ignored in this case. |

# API Endpoints

The following API endpoints are available:

- `/kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods`
  - `DELETE`: delete collection of Pod
  - `GET`: list objects of kind Pod
  - `POST`: create a new Pod
- `/kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods/{name}`
  - `DELETE`: delete the specified Pod
  - `GET`: read the specified Pod
  - `PATCH`: partially update the specified Pod
  - `PUT`: replace the specified Pod
- `/kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods/{name}/status`
  - `GET`: read status of the specified Pod
  - `PATCH`: partially update status of the specified Pod
  - `PUT`: replace status of the specified Pod

## /kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods

**HTTP method**

`DELETE`

**Description**

delete collection of Pod

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`GET`

**Description**

list objects of kind Pod

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `PodList` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`POST`

**Description**

create a new Pod

**Query parameters**

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

**Body parameters**

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `Pod` schema | `application/json` formatted |

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Pod` schema |
| 201 - Created | `Pod` schema |
| 202 - Accepted | `Pod` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods/{name}

### HTTP method

`DELETE`

### Description

delete the specified Pod

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 202 - Accepted | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

### Description

read the specified Pod

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Pod` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update the specified Pod

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Pod` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace the specified Pod

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `Pod` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Pod` schema |
| 201 - Created | `Pod` schema |

| HTTP code | Response body |
| --- | --- |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods/{name}/status

### HTTP method

`GET`

### Description

read status of the specified Pod

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Pod` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update status of the specified Pod

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Pod` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace status of the specified Pod

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

**Body parameters**

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `Pod` schema | `application/json` formatted |

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Pod` schema |
| 201 - Created | `Pod` schema |
| 401 - Unauthorized | Empty |

# /kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods/{name}/eviction

## Common Parameters

- `cluster` (*in path*): `string` required
  The name of the kuberentes cluster to access.
- `dryRun` (*in query*): `string`
  When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
- `fieldManager` (*in query*): `string`
  fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint↗.
- `fieldValidation` (*in query*): `string`
  fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
- `name` (*in path*): `string` required
  name of the Eviction
- `namespace` (*in path*): `string` required

object name and auth scope, such as for teams and projects

- `pretty` (*in query*): `string`
  If 'true', then the output is pretty printed. Defaults to 'false' unless the user-agent indicates a browser or command-line HTTP tool (curl and wget).

### `post`

create eviction of a Pod

#### Request Body

Eviction

#### Response

- `200` Eviction: OK
- `201` Eviction: Created
- `202` Eviction: Accepted
- `401` : Unauthorized

## Eviction

Eviction evicts a pod from its node subject to certain policies and safety constraints. This is a subresource of Pod. A request to cause such an eviction is created by POSTing to .../pods//evictions.

- `apiVersion` : `string`
  APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
- `deleteOptions` : DeleteOptions
  DeleteOptions may be provided when deleting an API object.
- `kind` : `string`
  Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
- `metadata` : ObjectMeta
  ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.

## /kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods/{name}/log

### Common Parameters

- `cluster` (*in path*): `string` required
  The name of the kuberentes cluster to access.
- `container` (*in query*): `string`
  The container for which to stream logs. Defaults to only container if there is one container in the pod.
- `follow` (*in query*): `boolean`
  Follow the log stream of the pod. Defaults to false.
- `insecureSkipTLSVerifyBackend` (*in query*): `boolean`

insecureSkipTLSVerifyBackend indicates that the apiserver should not confirm the validity of the serving certificate of the backend it is connecting to. This will make the HTTPS connection between the apiserver and the backend insecure. This means the apiserver cannot verify the log data it is receiving came from the real kubelet. If the kubelet is configured to verify the apiserver's TLS credentials, it does not mean the connection to the real kubelet is vulnerable to a man in the middle attack (e.g. an attacker could not intercept the actual log data coming from the real kubelet).

- `limitBytes` (*in query*): `integer`
  If set, the number of bytes to read from the server before terminating the log output. This may not display a complete final line of logging, and may return slightly more or slightly less than the specified limit.
- `name` (*in path*): `string` required
  name of the Pod
- `namespace` (*in path*): `string` required
  object name and auth scope, such as for teams and projects
- `pretty` (*in query*): `string`
  If 'true', then the output is pretty printed. Defaults to 'false' unless the user-agent indicates a browser or command-line HTTP tool (curl and wget).
- `previous` (*in query*): `boolean`
  Return previous terminated container logs. Defaults to false.
- `sinceSeconds` (*in query*): `integer`
  A relative time in seconds before the current time from which to show logs. If this value precedes the time a pod was started, only logs since the pod start will be returned. If this value is in the future, no logs will be returned. Only one of sinceSeconds or sinceTime may be specified.
- `stream` (*in query*): `string`
  Specify which container log stream to return to the client. Acceptable values are "All", "Stdout" and "Stderr". If not specified, "All" is used, and both stdout and stderr are returned interleaved. Note that when "TailLines" is specified, "Stream" can only be set to nil or "All".
- `tailLines` (*in query*): `integer`
  If set, the number of lines from the end of the logs to show. If not specified, logs are shown from the creation of the container or sinceSeconds or sinceTime. Note that when "TailLines" is specified, "Stream" can only be set to nil or "All".
- `timestamps` (*in query*): `boolean`
  If true, add an RFC3339 or RFC3339Nano timestamp at the beginning of every line of log output. Defaults to false.

### `get`

read log of the specified Pod

#### Response

- `200` `string` : OK
- `401` : Unauthorized

## /kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods/{name}/exec

## Common Parameters

- `cluster` (*in path*): `string` required
  The name of the kuberentes cluster to access.
- `command` (*in query*): `string`
  Command is the remote command to execute. argv array. Not executed within a shell.
- `container` (*in query*): `string`
  Container in which to execute the command. Defaults to only container if there is only one container in the pod.
- `name` (*in path*): `string` required

name of the PodExecOptions

- `namespace` (*in path*): `string` required
  object name and auth scope, such as for teams and projects
- `stderr` (*in query*): `boolean`
  Redirect the standard error stream of the pod for this call.
- `stdin` (*in query*): `boolean`
  Redirect the standard input stream of the pod for this call. Defaults to false.
- `stdout` (*in query*): `boolean`
  Redirect the standard output stream of the pod for this call.
- `tty` (*in query*): `boolean`
  TTY if true indicates that a tty will be allocated for the exec call. Defaults to false.

### `get`

connect GET requests to exec of Pod

#### Response

- `200` `string`: OK
- `401` : Unauthorized

### `post`

connect POST requests to exec of Pod

#### Response

- `200` `string`: OK
- `401` : Unauthorized

## /kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods/{name}/ephemeralcontainers

### Common Parameters

- `name` (*in path*): `string` required
  name of the Pod
- `namespace` (*in path*): `string` required
  object name and auth scope, such as for teams and projects
- `pretty` (*in query*): `string`
  If 'true', then the output is pretty printed. Defaults to 'false' unless the user-agent indicates a browser or command-line HTTP tool (curl and wget).

### `get`

read ephemeralcontainers of the specified Pod

#### Response

- `200` Pod: OK

- `401` : Unauthorized

## `put`

replace ephemeralcontainers of the specified Pod

### Parameters

- `dryRun` (*in query*): `string`
  When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
- `fieldManager` (*in query*): `string`
  fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint ↗.
- `fieldValidation` (*in query*): `string`
  fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

### Request Body

Pod

### Response

- `200` Pod: OK
- `201` Pod: Created
- `401` : Unauthorized

## `patch`

partially update ephemeralcontainers of the specified Pod

### Parameters

- `dryRun` (*in query*): `string`
  When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
- `fieldManager` (*in query*): `string`
  fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint ↗. This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
- `fieldValidation` (*in query*): `string`
  fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

- `force` (*in query*): `boolean`
  Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

### Request Body

Patch

### Response

- `200` Pod: OK
- `201` Pod: Created
- `401` : Unauthorized

## Patch

Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

# /kubernetes/{cluster}/api/v1/namespaces/{namespace}/pods/{name}/attach

## Common Parameters

- `container` (*in query*): `string`
  The container in which to execute the command. Defaults to only container if there is only one container in the pod.
- `name` (*in path*): `string` required
  name of the PodAttachOptions
- `namespace` (*in path*): `string` required
  object name and auth scope, such as for teams and projects
- `stderr` (*in query*): `boolean`
  Stderr if true indicates that stderr is to be redirected for the attach call. Defaults to true.
- `stdin` (*in query*): `boolean`
  Stdin if true, redirects the standard input stream of the pod for this call. Defaults to false.
- `stdout` (*in query*): `boolean`
  Stdout if true indicates that stdout is to be redirected for the attach call. Defaults to true.
- `tty` (*in query*): `boolean`
  TTY if true indicates that a tty will be allocated for the attach call. This is passed through the container runtime so the tty is allocated on the worker node by the container runtime. Defaults to false.

### `get`

connect GET requests to attach of Pod

### Response

- `200` `string`: OK
- `401` : Unauthorized

### `post`

connect POST requests to attach of Pod

### Response

- `200` `string` : OK
- `401` : Unauthorized

# Replicaset [apps/v1]

**Description**

ReplicaSet ensures that a specified number of pod replicas are running at any given time.

**Type**

`object`

## Specification

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources |
| `kind` | `string` | Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | ReplicaSetSpec is the specification of a ReplicaSet. |
| `status` | `object` | ReplicaSetStatus represents the current status of a ReplicaSet. |

### .spec

**Description**

ReplicaSetSpec is the specification of a ReplicaSet.

**Type**

`object`

**Required**

`selector`

| Property | Type | Description |
| --- | --- | --- |
| `minReadySeconds` | `integer` | Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready) |
| `replicas` | `integer` | Replicas is the number of desired replicas. This is a pointer to distinguish between explicit zero and unspecified. Defaults to 1. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller/#what-is-a-replicationcontroller ↗ |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `template` | `object` | PodTemplateSpec describes the data a pod should have when created from a template |

## .spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template

**Description**

PodTemplateSpec describes the data a pod should have when created from a template

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |

| Property | Type | Description |
| --- | --- | --- |
| `spec` | `object` | PodSpec is a description of a pod. |

## .spec.template.spec

### Description

PodSpec is a description of a pod.

### Type

`object`

### Required

`containers`

| Property | Type | Description |
| --- | --- | --- |
| `activeDeadlineSeconds` | `integer` | Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer. |
| `affinity` | `object` | Affinity is a group of affinity scheduling rules. |
| `automountServiceAccountToken` | `boolean` | AutomountServiceAccountToken indicates whether a service account token should be automatically mounted. |
| `containers` | `array` | List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated. |
| `dnsConfig` | `object` | PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy. |
| `dnsPolicy` | `string` | Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.<br><br>Possible enum values:<br>- `"ClusterFirst"` indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>- `"ClusterFirstWithHostNet"` indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>- `"Default"` indicates that the pod should use the default (as determined by kubelet) DNS settings. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"None"` indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig. |
| `enableServiceLinks` | `boolean` | EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching traditional container linking syntax. Optional: Defaults to true. |
| `ephemeralContainers` | `array` | List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. |
| `hostAliases` | `array` | HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. |
| `hostIPC` | `boolean` | Use the host's ipc namespace. Optional: Default to false. |
| `hostNetwork` | `boolean` | Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false. |
| `hostPID` | `boolean` | Use the host's pid namespace. Optional: Default to false. |
| `hostUsers` | `boolean` | Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new userns is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature. |
| `hostname` | `string` | Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value. |
| `imagePullSecrets` | `array` | ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod ↗ |
| `initContainers` | `array` | List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must |

| Property | Type | Description |
| --- | --- | --- |
| | | be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/ |
| `nodeName` | `string` | NodeName indicates in which node this pod is scheduled. If empty, this pod is a candidate for scheduling by the scheduler defined in schedulerName. Once this field is set, the kubelet for this node becomes responsible for the lifecycle of this pod. This field should not be used to express a desire for the pod to be scheduled on a specific node. https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodename |
| `nodeSelector` | `object` | NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/ |
| `os` | `object` | PodOS defines the OS parameters of a pod. |
| `overhead` | `object` | Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md |
| `preemptionPolicy` | `string` | PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.<br><br>Possible enum values:<br>- `"Never"` means that pod never preempts other pods with lower priority.<br>- `"PreemptLowerPriority"` means that pod can preempt other pods with lower priority. |
| `priority` | `integer` | The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority. |
| `priorityClassName` | `string` | If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest |

| Property | Type | Description |
| --- | --- | --- |
| | | priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default. |
| `readinessGates` | `array` | If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates ↗ |
| `resourceClaims` | `array` | ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.<br><br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy ↗<br><br>Possible enum values:<br>• `"Always"`<br>• `"Never"`<br>• `"OnFailure"` |
| `runtimeClassName` | `string` | RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class ↗ |
| `schedulerName` | `string` | If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler. |
| `schedulingGates` | `array` | SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.<br><br>SchedulingGates can only be set at pod creation time, and be removed only afterwards. |
| `securityContext` | `object` | PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext |

| Property | Type | Description |
| --- | --- | --- |
| | | take precedence over field values of PodSecurityContext. |
| `serviceAccount` | `string` | DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead. |
| `serviceAccountName` | `string` | ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/ ↗ |
| `setHostnameAsFQDN` | `boolean` | If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false. |
| `shareProcessNamespace` | `boolean` | Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false. |
| `subdomain` | `string` | If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all. |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds. |
| `tolerations` | `array` | If specified, the pod's tolerations. |
| `topologySpreadConstraints` | `array` | TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed. |
| `volumes` | `array` | List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes ↗ |

## .spec.template.spec.affinity

**Description**

Affinity is a group of affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `nodeAffinity` | `object` | Node affinity is a group of node affinity scheduling rules. |
| `podAffinity` | `object` | Pod affinity is a group of inter pod affinity scheduling rules. |
| `podAntiAffinity` | `object` | Pod anti affinity is a group of inter pod anti affinity scheduling rules. |

## .spec.template.spec.affinity.nodeAffinity

**Description**

Node affinity is a group of node affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `object` | A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

**Type**

`object`

**Required**

`weight` `preference`

| Property | Type | Description |
|---|---|---|
| `preference` | `object` | A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm. |
| `weight` | `integer` | Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

**Type**

`object`

**Required**

`nodeSelectorTerms`

| Property | Type | Description |
| --- | --- | --- |
| `nodeSelectorTerms` | `array` | Required. A list of node selector terms. The terms are ORed. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms

**Description**

Required. A list of node selector terms. The terms are ORed.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |

| Property | Type | Description |
| --- | --- | --- |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |

| Property | Type | Description |
| --- | --- | --- |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity

**Description**

Pod affinity is a group of inter pod affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |

| Property | Type | Description |
| --- | --- | --- |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
| --- | --- | --- |
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also,

mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and |

| Property | Type | Description |
| --- | --- | --- |
| | | labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity

**Description**

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the |

| Property | Type | Description |
| --- | --- | --- |
| | | elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
| --- | --- | --- |
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

# .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into |

| Property | Type | Description |
| --- | --- | --- |
| | | consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.containers

**Description**

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[]

**Description**

A single application container that you want to run within a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |

| Property | Type | Description |
| --- | --- | --- |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |

| Property | Type | Description |
| --- | --- | --- |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.containers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].args[]

**Type**

`string`

## .spec.template.spec.containers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].command[]

**Type**

`string`

## .spec.template.spec.containers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.containers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.containers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.containers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.containers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.containers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.containers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.containers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.containers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.containers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.containers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.containers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.containers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.containers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.containers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.containers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.containers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.containers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.containers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.containers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.containers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.containers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.containers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.containers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.containers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.containers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.dnsConfig

**Description**

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `nameservers` | `array` | A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed. |
| `options` | `array` | A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy. |
| `searches` | `array` | A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed. |

## .spec.template.spec.dnsConfig.nameservers

**Description**

A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.nameservers[]

**Type**

`string`

## .spec.template.spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

`array`

## .spec.template.spec.dnsConfig.options[]

**Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is this DNS resolver option's name. Required. |
| `value` | `string` | Value is this DNS resolver option's value. |

## .spec.template.spec.dnsConfig.searches

**Description**

A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.searches[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers

**Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[]

**Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable |

| Property | Type | Description |
| --- | --- | --- |
| | | exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | Ports are not allowed for ephemeral containers. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for the container to manage the restart behavior of each container within a pod. This may only be set for init containers. You cannot set this field on ephemeral containers. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `targetContainerName` | `string` | If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.<br><br>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined. |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message |

| Property | Type | Description |
| --- | --- | --- |
| | | length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>- `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>- `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.ephemeralContainers[].args

**Description**

Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].args[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.ephemeralContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.ephemeralContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].ports

**Description**

Ports are not allowed for ephemeral containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |

| Property | Type | Description |
| --- | --- | --- |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br><br>Possible enum values:<br>- `"SCTP"` is the SCTP protocol.<br>- `"TCP"` is the TCP protocol.<br>- `"UDP"` is the UDP protocol. |

## .spec.template.spec.ephemeralContainers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.ephemeralContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.ephemeralContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.ephemeralContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.ephemeralContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.ephemeralContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.ephemeralContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.ephemeralContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.ephemeralContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.hostAliases

**Description**

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified.

**Type**

`array`

## .spec.template.spec.hostAliases[]

**Description**

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

**Type**

`object`

**Required**

`ip`

| Property | Type | Description |
|---|---|---|
| `hostnames` | `array` | Hostnames for the above IP address. |
| `ip` | `string` | IP address of the host file entry. |

## .spec.template.spec.hostAliases[].hostnames

**Description**

Hostnames for the above IP address.

**Type**

`array`

## .spec.template.spec.hostAliases[].hostnames[]

**Type**

`string`

## .spec.template.spec.imagePullSecrets

**Description**

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

**Type**

`array`

## .spec.template.spec.imagePullSecrets[]

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.initContainers

### Description

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

### Type

`array`

## .spec.template.spec.initContainers[]

### Description

A single application container that you want to run within a pod.

### Type

`object`

### Required

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |

| Property | Type | Description |
| --- | --- | --- |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.initContainers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].args[]

**Type**

`string`

## .spec.template.spec.initContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].command[]

**Type**

`string`

## .spec.template.spec.initContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

# .spec.template.spec.initContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

# .spec.template.spec.initContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

# .spec.template.spec.initContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.initContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
|---|---|---|
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.initContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.initContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.initContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.initContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.initContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.initContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.initContainers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.initContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.initContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.initContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.initContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.initContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.initContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.initContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.initContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.initContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.initContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.initContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.initContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.initContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.nodeSelector

**Description**

NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

**Type**

`object`

## .spec.template.spec.os

**Description**

PodOS defines the OS parameters of a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration ↗ Clients should expect to handle additional values and treat unrecognized values in this field as os: null |

## .spec.template.spec.overhead

**Description**

Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md

**Type**

`object`

## .spec.template.spec.readinessGates

**Description**

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates

**Type**

`array`

## .spec.template.spec.readinessGates[]

**Description**

PodReadinessGate contains the reference to a pod condition

**Type**

`object`

**Required**

`conditionType`

| Property | Type | Description |
|---|---|---|
| `conditionType` | `string` | ConditionType refers to a condition in the pod's condition list with matching type. |

## .spec.template.spec.resourceClaims

**Description**

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable.

**Type**

`array`

## .spec.template.spec.resourceClaims[]

**Description**

PodResourceClaim references exactly one ResourceClaim, either directly or by naming a ResourceClaimTemplate which is then turned into a ResourceClaim for the pod. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL. |
| `resourceClaimName` | `string` | ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |
| `resourceClaimTemplateName` | `string` | ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.<br><br>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The pod name and resource name, along with a generated component, will be used to form a unique name for the ResourceClaim, which will be recorded in pod.status.resourceClaimStatuses.<br><br>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |

## .spec.template.spec.resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.template.spec.resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |

| Property | Type | Description |
| --- | --- | --- |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.schedulingGates

**Description**

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod. SchedulingGates can only be set at pod creation time, and be removed only afterwards.

**Type**

`array`

## .spec.template.spec.schedulingGates[]

**Description**

PodSchedulingGate is associated to a Pod to guard its scheduling.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the scheduling gate. Each scheduling gate must have a unique name field. |

## .spec.template.spec.securityContext

**Description**

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `fsGroup` | `integer` | A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:<br>1. The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw----<br>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows. |
| `fsGroupChangePolicy` | `string` | fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership(and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.<br>Possible enum values:<br>- `"Always"` indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior.<br>- `"OnRootMismatch"` indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `seLinuxChangePolicy` | `string` | seLinuxChangePolicy defines how the container's SELinux label is applied to all volumes used by the Pod. It has no effect on nodes that do not support SELinux or to volumes does not support SELinux. Valid values are "MountOption" and "Recursive".<br><br>"Recursive" means relabeling of all files on all Pod volumes by the container runtime. This may be slow for large volumes, but allows mixing privileged and unprivileged Pods sharing the same volume on the same node.<br><br>"MountOption" mounts all eligible Pod volumes with `-o context` mount option. This requires all Pods that share the same volume to use the same SELinux label. It is not possible to share the same volume among privileged and unprivileged Pods. Eligible volumes are in-tree FibreChannel and iSCSI volumes, and all CSI volumes whose CSI driver announces SELinux support by setting spec.seLinuxMount: true in their CSIDriver instance. Other volumes are always re-labelled recursively. "MountOption" value is allowed only when SELinuxMount feature gate is enabled.<br><br>If not specified and SELinuxMount feature gate is enabled, "MountOption" is used. If not specified and SELinuxMount feature gate is disabled, "MountOption" is used for ReadWriteOncePod volumes and "Recursive" for all other volumes.<br><br>This field affects only Pods that have SELinux label set, either in PodSecurityContext or in SecurityContext of all containers.<br><br>All Pods that use the same volume should use the same seLinuxChangePolicy, otherwise some pods can get stuck in ContainerCreating state. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `supplementalGroups` | `array` | A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows. |
| `supplementalGroupsPolicy` | `string` | Defines how supplemental groups of the first container processes are calculated. Valid values are "Merge" and "Strict". If not specified, "Merge" is used. (Alpha) Using the field requires the SupplementalGroupsPolicy feature gate to be enabled and the container runtime must implement support for this feature. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br><br>- `"Merge"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be merged with the primary user's groups as defined in the container image (in /etc/group). |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Strict"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be used instead of any groups defined in the container image. |
| `sysctls` | `array` | Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br>• `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>• `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `level` | `string` | Level is SELinux level label that applies to the container. |

| Property | Type | Description |
| --- | --- | --- |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.securityContext.supplementalGroups

**Description**

A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.supplementalGroups[]

**Type**

`integer`

## .spec.template.spec.securityContext.sysctls

**Description**

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.sysctls[]

**Description**

Sysctl defines a kernel parameter to be set

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of a property to set |
| `value` | `string` | Value of a property to set |

## .spec.template.spec.securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |

| Property | Type | Description |
| --- | --- | --- |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.tolerations

**Description**

If specified, the pod's tolerations.

**Type**

`array`

## .spec.template.spec.tolerations[]

**Description**

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `effect` | `string` | Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.<br><br>Possible enum values:<br>- `"NoExecute"` Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController.<br>- `"NoSchedule"` Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler.<br>- `"PreferNoSchedule"` Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler. |
| `key` | `string` | Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys. |
| `operator` | `string` | Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"Equal"`<br>- `"Exists"` |
| `tolerationSeconds` | `integer` | TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system. |
| `value` | `string` | Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string. |

## .spec.template.spec.topologySpreadConstraints

**Description**

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[]

**Description**

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

**Type**

`object`

**Required**

`maxSkew` `topologyKey` `whenUnsatisfiable`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.<br><br>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `maxSkew` | `integer` | MaxSkew describes the degree to which pods may be unevenly distributed. When `whenUnsatisfiable=DoNotSchedule`, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1. \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P \| - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When `whenUnsatisfiable=ScheduleAnyway`, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed. |
| `minDomains` | `integer` | MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.<br><br>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2: \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P P \| The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew. |
| `nodeAffinityPolicy` | `string` | NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.<br><br>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |
| `nodeTaintsPolicy` | `string` | NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.<br><br>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |

| Property | Type | Description |
| --- | --- | --- |
| `topologyKey` | `string` | TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field. |
| `whenUnsatisfiable` | `string` | WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: \| zone1 \| zone2 \| zone3 \| \| P P P \| P \| P \| If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it *more* imbalanced. It's a required field.<br><br>Possible enum values:<br>- `"DoNotSchedule"` instructs the scheduler not to schedule the pod when constraints are not satisfied.<br>- `"ScheduleAnyway"` instructs the scheduler to schedule the pod even if constraints are not satisfied. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector. This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.volumes

**Description**

List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

**Type**

`array`

## .spec.template.spec.volumes[]

**Description**

Volume represents a named volume in a pod that may be accessed by any container in the pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `awsElasticBlockStore` | `object` | Represents a Persistent Disk resource in AWS.<br><br>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling. |
| `azureDisk` | `object` | AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod. |
| `azureFile` | `object` | AzureFile represents an Azure File Service mount on the host and bind mount to the pod. |
| `cephfs` | `object` | Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling. |
| `cinder` | `object` | Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership |

| Property | Type | Description |
| --- | --- | --- |
| | | management and SELinux relabeling. |
| `configMap` | `object` | Adapts a ConfigMap into a volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling. |
| `csi` | `object` | Represents a source location of a volume to mount, managed by an external CSI driver |
| `downwardAPI` | `object` | DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling. |
| `emptyDir` | `object` | Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling. |
| `ephemeral` | `object` | Represents an ephemeral volume that is handled by a normal storage driver. |
| `fc` | `object` | Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling. |
| `flexVolume` | `object` | FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin. |
| `flocker` | `object` | Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling. |
| `gcePersistentDisk` | `object` | Represents a Persistent Disk resource in Google Compute Engine.<br><br>A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling. |
| `gitRepo` | `object` | Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.<br><br>DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container. |

| Property | Type | Description |
| --- | --- | --- |
| `glusterfs` | `object` | Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling. |
| `hostPath` | `object` | Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling. |
| `image` | `object` | ImageVolumeSource represents a image volume resource. |
| `iscsi` | `object` | Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling. |
| `name` | `string` | name of the volume. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `nfs` | `object` | Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling. |
| `persistentVolumeClaim` | `object` | PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system). |
| `photonPersistentDisk` | `object` | Represents a Photon Controller persistent disk resource. |
| `portworxVolume` | `object` | PortworxVolumeSource represents a Portworx volume resource. |
| `projected` | `object` | Represents a projected volume source |
| `quobyte` | `object` | Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling. |
| `rbd` | `object` | Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling. |
| `scaleIO` | `object` | ScaleIOVolumeSource represents a persistent ScaleIO volume |

| Property | Type | Description |
| --- | --- | --- |
| `secret` | `object` | Adapts a Secret into a volume.<br><br>The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling. |
| `storageos` | `object` | Represents a StorageOS persistent volume resource. |
| `vsphereVolume` | `object` | Represents a vSphere volume resource. |

## .spec.template.spec.volumes[].awsElasticBlockStore

**Description**

Represents a Persistent Disk resource in AWS. An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). |
| `readOnly` | `boolean` | readOnly value true will force the readOnly setting in VolumeMounts. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `volumeID` | `string` | volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |

## .spec.template.spec.volumes[].azureDisk

**Description**

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`diskName` `diskURI`

| Property | Type | Description |
|---|---|---|
| `cachingMode` | `string` | cachingMode is the Host Caching mode: None, Read Only, Read Write.<br>Possible enum values:<br>• `"None"`<br>• `"ReadOnly"`<br>• `"ReadWrite"` |
| `diskName` | `string` | diskName is the Name of the data disk in the blob storage |
| `diskURI` | `string` | diskURI is the URI of data disk in the blob storage |
| `fsType` | `string` | fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `kind` | `string` | kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared<br>Possible enum values:<br>• `"Dedicated"`<br>• `"Managed"`<br>• `"Shared"` |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

## .spec.template.spec.volumes[].azureFile

**Description**

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`secretName` `shareName`

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretName` | `string` | secretName is the name of secret that contains Azure Storage Account Name and Key |
| `shareName` | `string` | shareName is the azure share Name |

## .spec.template.spec.volumes[].cephfs

**Description**

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`monitors`

| Property | Type | Description |
| --- | --- | --- |
| `monitors` | `array` | monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `path` | `string` | path is Optional: Used as the mounted root, rather than the full Ceph tree, default is / |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretFile` | `string` | secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].cephfs.monitors

**Description**

monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].cephfs.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].cephfs.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].cinder

**Description**

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeID` | `string` | volumeID used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |

## .spec.template.spec.volumes[].cinder.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].configMap

**Description**

Adapts a ConfigMap into a volume. The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `defaultMode` | `integer` | defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].csi

**Description**

Represents a source location of a volume to mount, managed by an external CSI driver

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
|---|---|---|
| `driver` | `string` | driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster. |
| `fsType` | `string` | fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply. |

| Property | Type | Description |
| --- | --- | --- |
| `nodePublishSecretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `readOnly` | `boolean` | readOnly specifies a read-only configuration for the volume. Defaults to false (read/write). |
| `volumeAttributes` | `object` | volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values. |

## .spec.template.spec.volumes[].csi.nodePublishSecretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].csi.volumeAttributes

**Description**

volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**Type**

`object`

## .spec.template.spec.volumes[].downwardAPI

**Description**

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | Items is a list of downward API volume file |

## .spec.template.spec.volumes[].downwardAPI.items

**Description**

Items is a list of downward API volume file

**Type**

`array`

## .spec.template.spec.volumes[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
|---|---|---|
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].emptyDir

### Description

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `medium` | `string` | medium represents what type of storage medium should back this directory. The default is "" which means to use the n |
| `sizeLimit` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAM<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits> <r<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.html)<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a number<br>When a Quantity is parsed from a string, it will remember the type of suffix it had, an<br>Before serializing, Quantity will be put in "canonical form". This means that Exponent/<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suffix<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point number.<br>Non-canonical values will still parse as long as they are well formed, but will be re-e<br>This format is intended to make it difficult to use these numbers without writing some |

## .spec.template.spec.volumes[].ephemeral

**Description**

Represents an ephemeral volume that is handled by a normal storage driver.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `volumeClaimTemplate` | `object` | PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate

**Description**

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

**Type**

`object`

**Required**

`spec`

| Property | Type | Description |
|---|---|---|
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec

### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

### Type

`object`

| Property | Type | Description |
|---|---|---|
| `accessModes` | `array` | accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1 ↗ |
| `dataSource` | `object` | TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace. |
| `dataSourceRef` | `object` | TypedObjectReference contains enough information to let you locate the typed referenced object |
| `resources` | `object` | VolumeResourceRequirements describes the storage resource requirements for a volume. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `storageClassName` | `string` | storageClassName is the name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1 ↗ |
| `volumeAttributesClassName` | `string` | volumeAttributesClassName may be used to set the VolumeAttributesClass used by this claim. If specified, the CSI driver will create or update the volume with the attributes defined in the corresponding VolumeAttributesClass. This has a different purpose than storageClassName, it can be changed after the claim is created. An empty string value means that no VolumeAttributesClass will be applied to the claim but it's not allowed to reset this field to empty string once it is set. If unspecified and the PersistentVolumeClaim is unbound, the default VolumeAttributesClass will be set by the persistentvolume controller if it exists. If the resource referred to by volumeAttributesClass does not exist, this PersistentVolumeClaim will be set to a Pending state, as reflected by the |

| Property | Type | Description |
| --- | --- | --- |
| | | modifyVolumeStatus field, until such as a resource exists. More info: https://kubernetes.io/docs/concepts/storage/volume-attributes-classes/ (Beta) Using this field requires the VolumeAttributesClass feature gate to be enabled (off by default). |
| `volumeMode` | `string` | volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.<br>Possible enum values:<br>• `"Block"` means the volume will not be formatted with a filesystem and will remain a raw block device.<br>• `"Filesystem"` means the volume will be or is formatted with a filesystem. |
| `volumeName` | `string` | volumeName is the binding reference to the PersistentVolume backing this claim. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes

**Description**

accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

**Description**

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is the name of resource being referenced |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef

**Description**

TypedObjectReference contains enough information to let you locate the typed referenced object

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |
| `namespace` | `string` | Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

**Description**

VolumeResourceRequirements describes the storage resource requirements for a volume.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].fc

**Description**

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `lun` | `integer` | lun is Optional: FC target lun number |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `targetWWNs` | `array` | targetWWNs is Optional: FC target worldwide names (WWNs) |
| `wwids` | `array` | wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously. |

## .spec.template.spec.volumes[].fc.targetWWNs

**Description**

targetWWNs is Optional: FC target worldwide names (WWNs)

**Type**

`array`

## .spec.template.spec.volumes[].fc.targetWWNs[]

**Type**

`string`

## .spec.template.spec.volumes[].fc.wwids

**Description**

wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

**Type**

`array`

## .spec.template.spec.volumes[].fc.wwids[]

**Type**

`string`

## .spec.template.spec.volumes[].flexVolume

**Description**

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
| --- | --- | --- |
| `driver` | `string` | driver is the name of the driver to use for this volume. |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script. |
| `options` | `object` | options is Optional: this field holds extra command options if any. |
| `readOnly` | `boolean` | readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |

## .spec.template.spec.volumes[].flexVolume.options

**Description**

options is Optional: this field holds extra command options if any.

**Type**

`object`

## .spec.template.spec.volumes[].flexVolume.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: |

| Property | Type | Description |
| --- | --- | --- |
| | | https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].flocker

### Description

Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `datasetName` | `string` | datasetName is Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated |
| `datasetUUID` | `string` | datasetUUID is the UUID of the dataset. This is unique identifier of a Flocker dataset |

## .spec.template.spec.volumes[].gcePersistentDisk

### Description

Represents a Persistent Disk resource in Google Compute Engine. A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

### Type

`object`

### Required

`pdName`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `pdName` | `string` | pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

## .spec.template.spec.volumes[].gitRepo

**Description**

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

**Type**

`object`

**Required**

`repository`

| Property | Type | Description |
| --- | --- | --- |
| `directory` | `string` | directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name. |
| `repository` | `string` | repository is the URL |
| `revision` | `string` | revision is the commit hash for the specified revision. |

## .spec.template.spec.volumes[].glusterfs

**Description**

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`endpoints` `path`

| Property | Type | Description |
| --- | --- | --- |
| `endpoints` | `string` | endpoints is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |
| `path` | `string` | path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

## .spec.template.spec.volumes[].hostPath

### Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`path`

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗ |
| `type` | `string` | type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗<br><br>Possible enum values:<br>• `""` For backwards compatible, leave it empty if unset<br>• `"BlockDevice"` A block device must exist at the given path<br>• `"CharDevice"` A character device must exist at the given path<br>• `"Directory"` A directory must exist at the given path<br>• `"DirectoryOrCreate"` If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet.<br>• `"File"` A file must exist at the given path<br>• `"FileOrCreate"` If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet.<br>• `"Socket"` A UNIX socket must exist at the given path |

## .spec.template.spec.volumes[].image

### Description

ImageVolumeSource represents a image volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `pullPolicy` | `string` | Policy for pulling OCI objects. Possible values are: Always: the kubelet always attempts to pull the reference. Container creation will fail If the pull fails. Never: the kubelet never pulls the reference and only uses a local image or artifact. |

| Property | Type | Description |
| --- | --- | --- |
| | | Container creation will fail if the reference isn't present. IfNotPresent: the kubelet pulls if the reference isn't already present on disk. Container creation will fail if the reference isn't present and the pull fails. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise.<br><br>Possible enum values:<br><br>- `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>- `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>- `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `reference` | `string` | Required: Image or artifact reference to be used. Behaves in the same way as pod.spec.containers[*].image. Pull secrets will be assembled in the same way as for the container image by looking up node credentials, SA image pull secrets, and pod spec image pull secrets. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |

## .spec.template.spec.volumes[].iscsi

### Description

Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`targetPortal` `iqn` `lun`

| Property | Type | Description |
| --- | --- | --- |
| `chapAuthDiscovery` | `boolean` | chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication |
| `chapAuthSession` | `boolean` | chapAuthSession defines whether support iSCSI Session CHAP authentication |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi |
| `initiatorName` | `string` | initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection. |
| `iqn` | `string` | iqn is the target iSCSI Qualified Name. |

| Property | Type | Description |
| --- | --- | --- |
| `iscsiInterface` | `string` | iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp). |
| `lun` | `integer` | lun represents iSCSI Target Lun number. |
| `portals` | `array` | portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `targetPortal` | `string` | targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |

## .spec.template.spec.volumes[].iscsi.portals

**Description**

portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

**Type**

`array`

## .spec.template.spec.volumes[].iscsi.portals[]

**Type**

`string`

## .spec.template.spec.volumes[].iscsi.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].nfs

**Description**

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`server` `path`

| Property | Type | Description |
|---|---|---|
| `path` | `string` | path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |
| `readOnly` | `boolean` | readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |
| `server` | `string` | server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |

## .spec.template.spec.volumes[].persistentVolumeClaim

**Description**

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

**Type**

`object`

**Required**

`claimName`

| Property | Type | Description |
|---|---|---|
| `claimName` | `string` | claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims |
| `readOnly` | `boolean` | readOnly Will force the ReadOnly setting in VolumeMounts. Default false. |

## .spec.template.spec.volumes[].photonPersistentDisk

**Description**

Represents a Photon Controller persistent disk resource.

**Type**

`object`

**Required**

`pdID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `pdID` | `string` | pdID is the ID that identifies Photon Controller persistent disk |

## .spec.template.spec.volumes[].portworxVolume

**Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `volumeID` | `string` | volumeID uniquely identifies a Portworx volume |

## .spec.template.spec.volumes[].projected

**Description**

Represents a projected volume source

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `defaultMode` | `integer` | defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `sources` | `array` | sources is the list of volume projections. Each entry in this list handles one source. |

## .spec.template.spec.volumes[].projected.sources

**Description**

sources is the list of volume projections. Each entry in this list handles one source.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[]

**Description**

Projection that may be projected along with other supported volume types. Exactly one of these fields must be set.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `clusterTrustBundle` | `object` | ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem. |
| `configMap` | `object` | Adapts a ConfigMap into a projected volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode. |
| `downwardAPI` | `object` | Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode. |
| `secret` | `object` | Adapts a secret into a projected volume.<br><br>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode. |
| `serviceAccountToken` | `object` | ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise). |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle

**Description**

ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem.

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `name` | `string` | Select a single ClusterTrustBundle by object name. Mutually-exclusive with signerName and labelSelector. |
| `optional` | `boolean` | If true, don't block pod startup if the referenced ClusterTrustBundle(s) aren't available. If using name, then the named ClusterTrustBundle is allowed not to exist. If using signerName, then the combination of signerName and labelSelector is allowed to match zero ClusterTrustBundles. |
| `path` | `string` | Relative path from the volume root to write the bundle. |
| `signerName` | `string` | Select all ClusterTrustBundles that match this signer name. Mutually-exclusive with name. The contents of all selected ClusterTrustBundles will be unified and deduplicated. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].projected.sources[].configMap

**Description**

Adapts a ConfigMap into a projected volume. The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].projected.sources[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI

**Description**

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | Items is a list of DownwardAPIVolume file |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items

**Description**

Items is a list of DownwardAPIVolume file

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |

| Property | Type | Description |
| --- | --- | --- |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and |

| Property | Type | Description |
| --- | --- | --- |
| | | The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].projected.sources[].secret

### Description

Adapts a secret into a projected volume. The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional field specify whether the Secret or its key must be defined |

## .spec.template.spec.volumes[].projected.sources[].secret.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].serviceAccountToken

**Description**

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `audience` | `string` | audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver. |
| `expirationSeconds` | `integer` | expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours.Defaults to 1 hour and must be at least 10 minutes. |
| `path` | `string` | path is the path relative to the mount point of the file to project the token into. |

## .spec.template.spec.volumes[].quobyte

**Description**

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`registry` `volume`

| Property | Type | Description |
|---|---|---|
| `group` | `string` | group to map volume access to Default is no group |
| `readOnly` | `boolean` | readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false. |
| `registry` | `string` | registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes |
| `tenant` | `string` | tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin |
| `user` | `string` | user to map volume access to Defaults to serivceaccount user |

| Property | Type | Description |
| --- | --- | --- |
| `volume` | `string` | volume is a string that references an already created Quobyte volume by name. |

## .spec.template.spec.volumes[].rbd

### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`monitors` `image`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd ↗ |
| `image` | `string` | image is the rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `keyring` | `string` | keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `monitors` | `array` | monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `pool` | `string` | pool is the rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is the rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].rbd.monitors

**Description**

monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].rbd.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].rbd.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].scaleIO

**Description**

ScaleIOVolumeSource represents a persistent ScaleIO volume

**Type**

`object`

**Required**

`gateway` `system` `secretRef`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs". |
| `gateway` | `string` | gateway is the host address of the ScaleIO API Gateway. |
| `protectionDomain` | `string` | protectionDomain is the name of the ScaleIO Protection Domain for the configured storage. |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

| Property | Type | Description |
| --- | --- | --- |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `sslEnabled` | `boolean` | sslEnabled Flag enable/disable SSL communication with Gateway, default false |
| `storageMode` | `string` | storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned. |
| `storagePool` | `string` | storagePool is the ScaleIO Storage Pool associated with the protection domain. |
| `system` | `string` | system is the name of the storage system as configured in ScaleIO. |
| `volumeName` | `string` | volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source. |

## .spec.template.spec.volumes[].scaleIO.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].secret

**Description**

Adapts a Secret into a volume. The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON |

| Property | Type | Description |
| --- | --- | --- |
| | | requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `optional` | `boolean` | optional field specify whether the Secret or its keys must be defined |
| `secretName` | `string` | secretName is the name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret ↗ |

## .spec.template.spec.volumes[].secret.items

**Description**

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].storageos

**Description**

Represents a StorageOS persistent volume resource.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeName` | `string` | volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace. |
| `volumeNamespace` | `string` | volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created. |

## .spec.template.spec.volumes[].storageos.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].vsphereVolume

**Description**

Represents a vSphere volume resource.

**Type**

`object`

**Required**

`volumePath`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `storagePolicyID` | `string` | storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName. |
| `storagePolicyName` | `string` | storagePolicyName is the storage Policy Based Management (SPBM) profile name. |
| `volumePath` | `string` | volumePath is the path that identifies vSphere volume vmdk |

## .status

**Description**

ReplicaSetStatus represents the current status of a ReplicaSet.

**Type**

`object`

**Required**

`replicas`

| Property | Type | Description |
|---|---|---|
| `availableReplicas` | `integer` | The number of available replicas (ready for at least minReadySeconds) for this replica set. |
| `conditions` | `array` | Represents the latest available observations of a replica set's current state. |
| `fullyLabeledReplicas` | `integer` | The number of pods that have labels matching the labels of the pod template of the replicaset. |
| `observedGeneration` | `integer` | ObservedGeneration reflects the generation of the most recently observed ReplicaSet. |

| Property | Type | Description |
| --- | --- | --- |
| `readyReplicas` | `integer` | readyReplicas is the number of pods targeted by this ReplicaSet with a Ready Condition. |
| `replicas` | `integer` | Replicas is the most recently observed number of replicas. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller/#what-is-a-replicationcontroller ↗ |

## .status.conditions

**Description**

Represents the latest available observations of a replica set's current state.

**Type**

`array`

## .status.conditions[]

**Description**

ReplicaSetCondition describes the state of a replica set at a certain point.

**Type**

`object`

**Required**

`type` `status`

| Property | Type | Description |
| --- | --- | --- |
| `lastTransitionTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | A human readable message indicating details about the transition. |
| `reason` | `string` | The reason for the condition's last transition. |
| `status` | `string` | Status of the condition, one of True, False, Unknown. |
| `type` | `string` | Type of replica set condition. |

# API Endpoints

The following API endpoints are available:

- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/replicasets`
  - `DELETE` : delete collection of ReplicaSet
  - `GET` : list objects of kind ReplicaSet
  - `POST` : create a new ReplicaSet
- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/replicasets/{name}`
  - `DELETE` : delete the specified ReplicaSet
  - `GET` : read the specified ReplicaSet
  - `PATCH` : partially update the specified ReplicaSet
  - `PUT` : replace the specified ReplicaSet
- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/replicasets/{name}/status`
  - `GET` : read status of the specified ReplicaSet
  - `PATCH` : partially update status of the specified ReplicaSet
  - `PUT` : replace status of the specified ReplicaSet

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/replicasets

**HTTP method**

`DELETE`

**Description**

delete collection of ReplicaSet

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`GET`

**Description**

list objects of kind ReplicaSet

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicaSetList` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`POST`

**Description**

create a new ReplicaSet

**Query parameters**

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `ReplicaSet` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicaSet` schema |
| 201 - Created | `ReplicaSet` schema |
| 202 - Accepted | `ReplicaSet` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/replicasets/{name}

### HTTP method

`DELETE`

### Description

delete the specified ReplicaSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 202 - Accepted | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

**Description**

read the specified ReplicaSet

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicaSet` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`PATCH`

**Description**

partially update the specified ReplicaSet

**Query parameters**

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicaSet` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`PUT`

**Description**

replace the specified ReplicaSet

**Query parameters**

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default |

| Parameter | Type | Description |
| --- | --- | --- |
| | | behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `ReplicaSet` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicaSet` schema |
| 201 - Created | `ReplicaSet` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/replicasets/{name}/status

### HTTP method

`GET`

### Description

read status of the specified ReplicaSet

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicaSet` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update status of the specified ReplicaSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in |

| Parameter | Type | Description |
| --- | --- | --- |
| | | v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicaSet` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace status of the specified ReplicaSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `ReplicaSet` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicaSet` schema |
| 201 - Created | `ReplicaSet` schema |
| 401 - Unauthorized | Empty |

# ReplicationController [v1]

**Description**

ReplicationController represents the configuration of a replication controller.

**Type**

`object`

## Specification

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources ↗ |
| `kind` | `string` | Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds ↗ |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | ReplicationControllerSpec is the specification of a replication controller. |
| `status` | `object` | ReplicationControllerStatus represents the current status of a replication controller. |

### .spec

**Description**

ReplicationControllerSpec is the specification of a replication controller.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `minReadySeconds` | `integer` | Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready) |

| Property | Type | Description |
| --- | --- | --- |
| `replicas` | `integer` | Replicas is the number of desired replicas. This is a pointer to distinguish between explicit zero and unspecified. Defaults to 1. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#what-is-a-replicationcontroller ↗ |
| `selector` | `object` | Selector is a label query over pods that should match the Replicas count. If Selector is empty, it is defaulted to the labels present on the Pod template. Label keys and values that must match in order to be controlled by this replication controller, if empty defaulted to labels on Pod template. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors ↗ |
| `template` | `object` | PodTemplateSpec describes the data a pod should have when created from a template |

## .spec.selector

**Description**

Selector is a label query over pods that should match the Replicas count. If Selector is empty, it is defaulted to the labels present on the Pod template. Label keys and values that must match in order to be controlled by this replication controller, if empty defaulted to labels on Pod template. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

**Type**

`object`

## .spec.template

**Description**

PodTemplateSpec describes the data a pod should have when created from a template

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PodSpec is a description of a pod. |

## .spec.template.spec

**Description**

PodSpec is a description of a pod.

**Type**

`object`

**Required**

`containers`

| Property | Type | Description |
| --- | --- | --- |
| `activeDeadlineSeconds` | `integer` | Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer. |
| `affinity` | `object` | Affinity is a group of affinity scheduling rules. |
| `automountServiceAccountToken` | `boolean` | AutomountServiceAccountToken indicates whether a service account token should be automatically mounted. |
| `containers` | `array` | List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated. |
| `dnsConfig` | `object` | PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy. |
| `dnsPolicy` | `string` | Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.<br><br>Possible enum values:<br>• `"ClusterFirst"` indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>• `"ClusterFirstWithHostNet"` indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>• `"Default"` indicates that the pod should use the default (as determined by kubelet) DNS settings.<br>• `"None"` indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig. |
| `enableServiceLinks` | `boolean` | EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching traditional container linking syntax. Optional: Defaults to true. |
| `ephemeralContainers` | `array` | List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. |

| Property | Type | Description |
| --- | --- | --- |
| `hostAliases` | `array` | HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. |
| `hostIPC` | `boolean` | Use the host's ipc namespace. Optional: Default to false. |
| `hostNetwork` | `boolean` | Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false. |
| `hostPID` | `boolean` | Use the host's pid namespace. Optional: Default to false. |
| `hostUsers` | `boolean` | Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new userns is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature. |
| `hostname` | `string` | Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value. |
| `imagePullSecrets` | `array` | ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod ↗ |
| `initContainers` | `array` | List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/ ↗ |
| `nodeName` | `string` | NodeName indicates in which node this pod is scheduled. If empty, this pod is a candidate for scheduling by the scheduler defined in schedulerName. Once this field is set, the kubelet for this node becomes responsible for the lifecycle of this pod. This field should not be used to express a desire for the pod to be scheduled on a specific node. https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodename ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `nodeSelector` | `object` | NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/ |
| `os` | `object` | PodOS defines the OS parameters of a pod. |
| `overhead` | `object` | Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md |
| `preemptionPolicy` | `string` | PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.<br><br>Possible enum values:<br>• `"Never"` means that pod never preempts other pods with lower priority.<br>• `"PreemptLowerPriority"` means that pod can preempt other pods with lower priority. |
| `priority` | `integer` | The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority. |
| `priorityClassName` | `string` | If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default. |
| `readinessGates` | `array` | If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates |

| Property | Type | Description |
| --- | --- | --- |
| `resourceClaims` | `array` | ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.<br><br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy ↗<br><br>Possible enum values:<br>• `"Always"`<br>• `"Never"`<br>• `"OnFailure"` |
| `runtimeClassName` | `string` | RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class ↗ |
| `schedulerName` | `string` | If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler. |
| `schedulingGates` | `array` | SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.<br><br>SchedulingGates can only be set at pod creation time, and be removed only afterwards. |
| `securityContext` | `object` | PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext. |
| `serviceAccount` | `string` | DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead. |

| Property | Type | Description |
| --- | --- | --- |
| `serviceAccountName` | `string` | ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/ |
| `setHostnameAsFQDN` | `boolean` | If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false. |
| `shareProcessNamespace` | `boolean` | Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false. |
| `subdomain` | `string` | If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all. |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds. |
| `tolerations` | `array` | If specified, the pod's tolerations. |
| `topologySpreadConstraints` | `array` | TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed. |
| `volumes` | `array` | List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes |

## .spec.template.spec.affinity

### Description

Affinity is a group of affinity scheduling rules.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `nodeAffinity` | `object` | Node affinity is a group of node affinity scheduling rules. |
| `podAffinity` | `object` | Pod affinity is a group of inter pod affinity scheduling rules. |
| `podAntiAffinity` | `object` | Pod anti affinity is a group of inter pod anti affinity scheduling rules. |

## .spec.template.spec.affinity.nodeAffinity

**Description**

Node affinity is a group of node affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `object` | A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

**Type**

`object`

**Required**

`weight` `preference`

| Property | Type | Description |
|---|---|---|
| `preference` | `object` | A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm. |
| `weight` | `integer` | Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

**Type**

`object`

**Required**

`nodeSelectorTerms`

| Property | Type | Description |
|---|---|---|
| `nodeSelectorTerms` | `array` | Required. A list of node selector terms. The terms are ORed. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms

**Description**

Required. A list of node selector terms. The terms are ORed.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity

**Description**

Pod affinity is a group of inter pod affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
|---|---|---|
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` |

| Property | Type | Description |
|---|---|---|
| | | as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

### Description

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

### Type

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or

may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |

| Property | Type | Description |
| --- | --- | --- |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity

**Description**

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
|---|---|---|
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |

| Property | Type | Description |
| --- | --- | --- |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only |

| Property | Type | Description |
| --- | --- | --- |
| | | "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.containers

**Description**

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[]

**Description**

A single application container that you want to run within a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |

| Property | Type | Description |
| --- | --- | --- |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |

| Property | Type | Description |
| --- | --- | --- |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.containers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].args[]

**Type**

`string`

## .spec.template.spec.containers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].command[]

**Type**

`string`

## .spec.template.spec.containers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.containers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.containers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.containers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.containers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.containers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.containers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.containers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.containers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.containers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.containers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.postStart.exec.command[]

**Type**

`string`

# .spec.template.spec.containers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

# .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

# .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.containers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.containers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.containers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.containers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.containers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.containers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.containers[].securityContext.appArmorProfile

### Description

AppArmorProfile defines a pod or container's AppArmor settings.

### Type

`object`

### Required

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.containers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.containers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.containers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.containers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa ↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.containers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.containers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.containers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.containers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.dnsConfig

**Description**

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `nameservers` | `array` | A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed. |
| `options` | `array` | A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy. |
| `searches` | `array` | A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed. |

## .spec.template.spec.dnsConfig.nameservers

**Description**

A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.nameservers[]

**Type**

`string`

## .spec.template.spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

`array`

## .spec.template.spec.dnsConfig.options[]

**Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is this DNS resolver option's name. Required. |
| `value` | `string` | Value is this DNS resolver option's value. |

## .spec.template.spec.dnsConfig.searches

**Description**

A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.searches[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers

**Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[]

**Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable |

| Property | Type | Description |
| --- | --- | --- |
| | | exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | Ports are not allowed for ephemeral containers. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for the container to manage the restart behavior of each container within a pod. This may only be set for init containers. You cannot set this field on ephemeral containers. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `targetContainerName` | `string` | If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.<br><br>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined. |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message |

| Property | Type | Description |
| --- | --- | --- |
| | | length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>- `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>- `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.ephemeralContainers[].args

**Description**

Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].args[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.ephemeralContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.ephemeralContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].ports

**Description**

Ports are not allowed for ephemeral containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |

| Property | Type | Description |
| --- | --- | --- |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br><br>Possible enum values:<br>- `"SCTP"` is the SCTP protocol.<br>- `"TCP"` is the TCP protocol.<br>- `"UDP"` is the UDP protocol. |

## .spec.template.spec.ephemeralContainers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.ephemeralContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.ephemeralContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.ephemeralContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.ephemeralContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.ephemeralContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.ephemeralContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.ephemeralContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.ephemeralContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.hostAliases

**Description**

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified.

**Type**

`array`

## .spec.template.spec.hostAliases[]

**Description**

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

**Type**

`object`

**Required**

`ip`

| Property | Type | Description |
|---|---|---|
| `hostnames` | `array` | Hostnames for the above IP address. |
| `ip` | `string` | IP address of the host file entry. |

## .spec.template.spec.hostAliases[].hostnames

**Description**

Hostnames for the above IP address.

**Type**

`array`

## .spec.template.spec.hostAliases[].hostnames[]

**Type**

`string`

## .spec.template.spec.imagePullSecrets

**Description**

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

**Type**

`array`

## .spec.template.spec.imagePullSecrets[]

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.initContainers

### Description

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

### Type

`array`

## .spec.template.spec.initContainers[]

### Description

A single application container that you want to run within a pod.

### Type

`object`

### Required

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |

| Property | Type | Description |
| --- | --- | --- |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.initContainers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].args[]

**Type**

`string`

## .spec.template.spec.initContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].command[]

**Type**

`string`

## .spec.template.spec.initContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

# .spec.template.spec.initContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

# .spec.template.spec.initContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

# .spec.template.spec.initContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.initContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
|---|---|---|
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.initContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.initContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.initContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.initContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.initContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.initContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

# .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].livenessProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.initContainers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.initContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.initContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.initContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.initContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.initContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.initContainers[].securityContext.appArmorProfile

### Description

AppArmorProfile defines a pod or container's AppArmor settings.

### Type

`object`

### Required

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.initContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.initContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.initContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.initContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.initContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.initContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.initContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.nodeSelector

**Description**

NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

**Type**

`object`

## .spec.template.spec.os

**Description**

PodOS defines the OS parameters of a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration ↗ Clients should expect to handle additional values and treat unrecognized values in this field as os: null |

## .spec.template.spec.overhead

**Description**

Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md

**Type**

`object`

## .spec.template.spec.readinessGates

**Description**

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates

**Type**

`array`

## .spec.template.spec.readinessGates[]

**Description**

PodReadinessGate contains the reference to a pod condition

**Type**

`object`

**Required**

`conditionType`

| Property | Type | Description |
|---|---|---|
| `conditionType` | `string` | ConditionType refers to a condition in the pod's condition list with matching type. |

## .spec.template.spec.resourceClaims

**Description**

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable.

**Type**

`array`

## .spec.template.spec.resourceClaims[]

**Description**

PodResourceClaim references exactly one ResourceClaim, either directly or by naming a ResourceClaimTemplate which is then turned into a ResourceClaim for the pod. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL. |
| `resourceClaimName` | `string` | ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |
| `resourceClaimTemplateName` | `string` | ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.<br><br>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The pod name and resource name, along with a generated component, will be used to form a unique name for the ResourceClaim, which will be recorded in pod.status.resourceClaimStatuses.<br><br>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |

## .spec.template.spec.resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.template.spec.resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |

| Property | Type | Description |
| --- | --- | --- |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.schedulingGates

**Description**

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod. SchedulingGates can only be set at pod creation time, and be removed only afterwards.

**Type**

`array`

## .spec.template.spec.schedulingGates[]

**Description**

PodSchedulingGate is associated to a Pod to guard its scheduling.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the scheduling gate. Each scheduling gate must have a unique name field. |

## .spec.template.spec.securityContext

**Description**

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `fsGroup` | `integer` | A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:<br>1. The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw----<br>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows. |
| `fsGroupChangePolicy` | `string` | fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership(and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.<br>Possible enum values:<br>- `"Always"` indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior.<br>- `"OnRootMismatch"` indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `seLinuxChangePolicy` | `string` | seLinuxChangePolicy defines how the container's SELinux label is applied to all volumes used by the Pod. It has no effect on nodes that do not support SELinux or to volumes does not support SELinux. Valid values are "MountOption" and "Recursive".<br><br>"Recursive" means relabeling of all files on all Pod volumes by the container runtime. This may be slow for large volumes, but allows mixing privileged and unprivileged Pods sharing the same volume on the same node.<br><br>"MountOption" mounts all eligible Pod volumes with `-o context` mount option. This requires all Pods that share the same volume to use the same SELinux label. It is not possible to share the same volume among privileged and unprivileged Pods. Eligible volumes are in-tree FibreChannel and iSCSI volumes, and all CSI volumes whose CSI driver announces SELinux support by setting spec.seLinuxMount: true in their CSIDriver instance. Other volumes are always re-labelled recursively. "MountOption" value is allowed only when SELinuxMount feature gate is enabled.<br><br>If not specified and SELinuxMount feature gate is enabled, "MountOption" is used. If not specified and SELinuxMount feature gate is disabled, "MountOption" is used for ReadWriteOncePod volumes and "Recursive" for all other volumes.<br><br>This field affects only Pods that have SELinux label set, either in PodSecurityContext or in SecurityContext of all containers.<br><br>All Pods that use the same volume should use the same seLinuxChangePolicy, otherwise some pods can get stuck in ContainerCreating state. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `supplementalGroups` | `array` | A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows. |
| `supplementalGroupsPolicy` | `string` | Defines how supplemental groups of the first container processes are calculated. Valid values are "Merge" and "Strict". If not specified, "Merge" is used. (Alpha) Using the field requires the SupplementalGroupsPolicy feature gate to be enabled and the container runtime must implement support for this feature. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br><br>- `"Merge"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be merged with the primary user's groups as defined in the container image (in /etc/group). |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Strict"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be used instead of any groups defined in the container image. |
| `sysctls` | `array` | Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br>• `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>• `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `level` | `string` | Level is SELinux level label that applies to the container. |

| Property | Type | Description |
| --- | --- | --- |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br><br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.securityContext.supplementalGroups

**Description**

A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.supplementalGroups[]

**Type**

`integer`

## .spec.template.spec.securityContext.sysctls

**Description**

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.sysctls[]

**Description**

Sysctl defines a kernel parameter to be set

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of a property to set |
| `value` | `string` | Value of a property to set |

## .spec.template.spec.securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |

| Property | Type | Description |
| --- | --- | --- |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.tolerations

**Description**

If specified, the pod's tolerations.

**Type**

`array`

## .spec.template.spec.tolerations[]

**Description**

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `effect` | `string` | Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.<br><br>Possible enum values:<br>- `"NoExecute"` Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController.<br>- `"NoSchedule"` Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler.<br>- `"PreferNoSchedule"` Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler. |
| `key` | `string` | Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys. |
| `operator` | `string` | Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"Equal"`<br>- `"Exists"` |
| `tolerationSeconds` | `integer` | TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system. |
| `value` | `string` | Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string. |

## .spec.template.spec.topologySpreadConstraints

**Description**

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[]

**Description**

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

**Type**

`object`

**Required**

`maxSkew` `topologyKey` `whenUnsatisfiable`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.<br><br>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `maxSkew` | `integer` | MaxSkew describes the degree to which pods may be unevenly distributed. When `whenUnsatisfiable=DoNotSchedule`, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1. \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P \| - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When `whenUnsatisfiable=ScheduleAnyway`, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed. |
| `minDomains` | `integer` | MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.<br><br>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2: \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P P \| The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew. |
| `nodeAffinityPolicy` | `string` | NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.<br><br>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |
| `nodeTaintsPolicy` | `string` | NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.<br><br>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |

| Property | Type | Description |
| --- | --- | --- |
| `topologyKey` | `string` | TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field. |
| `whenUnsatisfiable` | `string` | WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: \| zone1 \| zone2 \| zone3 \| \| P P P \| P \| P \| If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it *more* imbalanced. It's a required field.<br><br>Possible enum values:<br>- `"DoNotSchedule"` instructs the scheduler not to schedule the pod when constraints are not satisfied.<br>- `"ScheduleAnyway"` instructs the scheduler to schedule the pod even if constraints are not satisfied. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector. This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.volumes

**Description**

List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

**Type**

`array`

## .spec.template.spec.volumes[]

**Description**

Volume represents a named volume in a pod that may be accessed by any container in the pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `awsElasticBlockStore` | `object` | Represents a Persistent Disk resource in AWS.<br><br>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling. |
| `azureDisk` | `object` | AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod. |
| `azureFile` | `object` | AzureFile represents an Azure File Service mount on the host and bind mount to the pod. |
| `cephfs` | `object` | Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling. |
| `cinder` | `object` | Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership |

| Property | Type | Description |
| --- | --- | --- |
| | | management and SELinux relabeling. |
| `configMap` | `object` | Adapts a ConfigMap into a volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling. |
| `csi` | `object` | Represents a source location of a volume to mount, managed by an external CSI driver |
| `downwardAPI` | `object` | DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling. |
| `emptyDir` | `object` | Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling. |
| `ephemeral` | `object` | Represents an ephemeral volume that is handled by a normal storage driver. |
| `fc` | `object` | Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling. |
| `flexVolume` | `object` | FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin. |
| `flocker` | `object` | Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling. |
| `gcePersistentDisk` | `object` | Represents a Persistent Disk resource in Google Compute Engine.<br><br>A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling. |
| `gitRepo` | `object` | Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.<br><br>DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container. |

| Property | Type | Description |
| --- | --- | --- |
| `glusterfs` | `object` | Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling. |
| `hostPath` | `object` | Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling. |
| `image` | `object` | ImageVolumeSource represents a image volume resource. |
| `iscsi` | `object` | Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling. |
| `name` | `string` | name of the volume. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `nfs` | `object` | Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling. |
| `persistentVolumeClaim` | `object` | PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system). |
| `photonPersistentDisk` | `object` | Represents a Photon Controller persistent disk resource. |
| `portworxVolume` | `object` | PortworxVolumeSource represents a Portworx volume resource. |
| `projected` | `object` | Represents a projected volume source |
| `quobyte` | `object` | Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling. |
| `rbd` | `object` | Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling. |
| `scaleIO` | `object` | ScaleIOVolumeSource represents a persistent ScaleIO volume |

| Property | Type | Description |
| --- | --- | --- |
| `secret` | `object` | Adapts a Secret into a volume.<br><br>The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling. |
| `storageos` | `object` | Represents a StorageOS persistent volume resource. |
| `vsphereVolume` | `object` | Represents a vSphere volume resource. |

## .spec.template.spec.volumes[].awsElasticBlockStore

**Description**

Represents a Persistent Disk resource in AWS. An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). |
| `readOnly` | `boolean` | readOnly value true will force the readOnly setting in VolumeMounts. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `volumeID` | `string` | volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |

## .spec.template.spec.volumes[].azureDisk

**Description**

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`diskName` `diskURI`

| Property | Type | Description |
|---|---|---|
| `cachingMode` | `string` | cachingMode is the Host Caching mode: None, Read Only, Read Write.<br><br>Possible enum values:<br>• `"None"`<br>• `"ReadOnly"`<br>• `"ReadWrite"` |
| `diskName` | `string` | diskName is the Name of the data disk in the blob storage |
| `diskURI` | `string` | diskURI is the URI of data disk in the blob storage |
| `fsType` | `string` | fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `kind` | `string` | kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared<br><br>Possible enum values:<br>• `"Dedicated"`<br>• `"Managed"`<br>• `"Shared"` |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

## .spec.template.spec.volumes[].azureFile

**Description**

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`secretName` `shareName`

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretName` | `string` | secretName is the name of secret that contains Azure Storage Account Name and Key |
| `shareName` | `string` | shareName is the azure share Name |

## .spec.template.spec.volumes[].cephfs

**Description**

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`monitors`

| Property | Type | Description |
| --- | --- | --- |
| `monitors` | `array` | monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `path` | `string` | path is Optional: Used as the mounted root, rather than the full Ceph tree, default is / |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretFile` | `string` | secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].cephfs.monitors

**Description**

monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].cephfs.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].cephfs.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].cinder

**Description**

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeID` | `string` | volumeID used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |

## .spec.template.spec.volumes[].cinder.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].configMap

**Description**

Adapts a ConfigMap into a volume. The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `defaultMode` | `integer` | defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].csi

**Description**

Represents a source location of a volume to mount, managed by an external CSI driver

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
|---|---|---|
| `driver` | `string` | driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster. |
| `fsType` | `string` | fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply. |

| Property | Type | Description |
| --- | --- | --- |
| `nodePublishSecretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `readOnly` | `boolean` | readOnly specifies a read-only configuration for the volume. Defaults to false (read/write). |
| `volumeAttributes` | `object` | volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values. |

## .spec.template.spec.volumes[].csi.nodePublishSecretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].csi.volumeAttributes

**Description**

volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**Type**

`object`

## .spec.template.spec.volumes[].downwardAPI

**Description**

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | Items is a list of downward API volume file |

## .spec.template.spec.volumes[].downwardAPI.items

**Description**

Items is a list of downward API volume file

**Type**

`array`

## .spec.template.spec.volumes[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
|---|---|---|
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].emptyDir

### Description

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `medium` | `string` | medium represents what type of storage medium should back this directory. The default is "" which means to use the n |
| `sizeLimit` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAM<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits> <r<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.html)<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a number<br>When a Quantity is parsed from a string, it will remember the type of suffix it had, an<br>Before serializing, Quantity will be put in "canonical form". This means that Exponent/<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suffix<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point number.<br>Non-canonical values will still parse as long as they are well formed, but will be re-e<br>This format is intended to make it difficult to use these numbers without writing some |

## .spec.template.spec.volumes[].ephemeral

**Description**

Represents an ephemeral volume that is handled by a normal storage driver.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `volumeClaimTemplate` | `object` | PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate

**Description**

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

**Type**

`object`

**Required**

`spec`

| Property | Type | Description |
|---|---|---|
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec

### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

### Type

`object`

| Property | Type | Description |
|---|---|---|
| `accessModes` | `array` | accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1 ↗ |
| `dataSource` | `object` | TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace. |
| `dataSourceRef` | `object` | TypedObjectReference contains enough information to let you locate the typed referenced object |
| `resources` | `object` | VolumeResourceRequirements describes the storage resource requirements for a volume. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `storageClassName` | `string` | storageClassName is the name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1 ↗ |
| `volumeAttributesClassName` | `string` | volumeAttributesClassName may be used to set the VolumeAttributesClass used by this claim. If specified, the CSI driver will create or update the volume with the attributes defined in the corresponding VolumeAttributesClass. This has a different purpose than storageClassName, it can be changed after the claim is created. An empty string value means that no VolumeAttributesClass will be applied to the claim but it's not allowed to reset this field to empty string once it is set. If unspecified and the PersistentVolumeClaim is unbound, the default VolumeAttributesClass will be set by the persistentvolume controller if it exists. If the resource referred to by volumeAttributesClass does not exist, this PersistentVolumeClaim will be set to a Pending state, as reflected by the |

| Property | Type | Description |
| --- | --- | --- |
| | | modifyVolumeStatus field, until such as a resource exists. More info: https://kubernetes.io/docs/concepts/storage/volume-attributes-classes/ (Beta) Using this field requires the VolumeAttributesClass feature gate to be enabled (off by default). |
| `volumeMode` | `string` | volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.<br><br>Possible enum values:<br>• `"Block"` means the volume will not be formatted with a filesystem and will remain a raw block device.<br>• `"Filesystem"` means the volume will be or is formatted with a filesystem. |
| `volumeName` | `string` | volumeName is the binding reference to the PersistentVolume backing this claim. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes

**Description**

accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

**Description**

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is the name of resource being referenced |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef

**Description**

TypedObjectReference contains enough information to let you locate the typed referenced object

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |
| `namespace` | `string` | Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

**Description**

VolumeResourceRequirements describes the storage resource requirements for a volume.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].fc

**Description**

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `lun` | `integer` | lun is Optional: FC target lun number |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `targetWWNs` | `array` | targetWWNs is Optional: FC target worldwide names (WWNs) |
| `wwids` | `array` | wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously. |

## .spec.template.spec.volumes[].fc.targetWWNs

**Description**

targetWWNs is Optional: FC target worldwide names (WWNs)

**Type**

`array`

## .spec.template.spec.volumes[].fc.targetWWNs[]

**Type**

`string`

## .spec.template.spec.volumes[].fc.wwids

**Description**

wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

**Type**

`array`

## .spec.template.spec.volumes[].fc.wwids[]

**Type**

`string`

## .spec.template.spec.volumes[].flexVolume

**Description**

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
|---|---|---|
| `driver` | `string` | driver is the name of the driver to use for this volume. |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script. |
| `options` | `object` | options is Optional: this field holds extra command options if any. |
| `readOnly` | `boolean` | readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |

## .spec.template.spec.volumes[].flexVolume.options

**Description**

options is Optional: this field holds extra command options if any.

**Type**

`object`

## .spec.template.spec.volumes[].flexVolume.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: |

| Property | Type | Description |
| --- | --- | --- |
| | | https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].flocker

### Description

Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `datasetName` | `string` | datasetName is Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated |
| `datasetUUID` | `string` | datasetUUID is the UUID of the dataset. This is unique identifier of a Flocker dataset |

## .spec.template.spec.volumes[].gcePersistentDisk

### Description

Represents a Persistent Disk resource in Google Compute Engine. A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

### Type

`object`

### Required

`pdName`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `pdName` | `string` | pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

## .spec.template.spec.volumes[].gitRepo

**Description**

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

**Type**

`object`

**Required**

`repository`

| Property | Type | Description |
| --- | --- | --- |
| `directory` | `string` | directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name. |
| `repository` | `string` | repository is the URL |
| `revision` | `string` | revision is the commit hash for the specified revision. |

## .spec.template.spec.volumes[].glusterfs

**Description**

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`endpoints` `path`

| Property | Type | Description |
| --- | --- | --- |
| `endpoints` | `string` | endpoints is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |
| `path` | `string` | path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

## .spec.template.spec.volumes[].hostPath

### Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`path`

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗ |
| `type` | `string` | type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗<br><br>Possible enum values:<br>• `""` For backwards compatible, leave it empty if unset<br>• `"BlockDevice"` A block device must exist at the given path<br>• `"CharDevice"` A character device must exist at the given path<br>• `"Directory"` A directory must exist at the given path<br>• `"DirectoryOrCreate"` If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet.<br>• `"File"` A file must exist at the given path<br>• `"FileOrCreate"` If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet.<br>• `"Socket"` A UNIX socket must exist at the given path |

## .spec.template.spec.volumes[].image

### Description

ImageVolumeSource represents a image volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `pullPolicy` | `string` | Policy for pulling OCI objects. Possible values are: Always: the kubelet always attempts to pull the reference. Container creation will fail If the pull fails. Never: the kubelet never pulls the reference and only uses a local image or artifact. |

| Property | Type | Description |
| --- | --- | --- |
| | | Container creation will fail if the reference isn't present. IfNotPresent: the kubelet pulls if the reference isn't already present on disk. Container creation will fail if the reference isn't present and the pull fails. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise.<br><br>Possible enum values:<br><br>- `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>- `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>- `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `reference` | `string` | Required: Image or artifact reference to be used. Behaves in the same way as pod.spec.containers[*].image. Pull secrets will be assembled in the same way as for the container image by looking up node credentials, SA image pull secrets, and pod spec image pull secrets. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |

## .spec.template.spec.volumes[].iscsi

### Description

Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`targetPortal` `iqn` `lun`

| Property | Type | Description |
| --- | --- | --- |
| `chapAuthDiscovery` | `boolean` | chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication |
| `chapAuthSession` | `boolean` | chapAuthSession defines whether support iSCSI Session CHAP authentication |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi |
| `initiatorName` | `string` | initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection. |
| `iqn` | `string` | iqn is the target iSCSI Qualified Name. |

| Property | Type | Description |
| --- | --- | --- |
| `iscsiInterface` | `string` | iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp). |
| `lun` | `integer` | lun represents iSCSI Target Lun number. |
| `portals` | `array` | portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `targetPortal` | `string` | targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |

## .spec.template.spec.volumes[].iscsi.portals

**Description**

portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

**Type**

`array`

## .spec.template.spec.volumes[].iscsi.portals[]

**Type**

`string`

## .spec.template.spec.volumes[].iscsi.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].nfs

**Description**

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`server` `path`

| Property | Type | Description |
|---|---|---|
| `path` | `string` | path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs ↗ |
| `readOnly` | `boolean` | readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs ↗ |
| `server` | `string` | server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs ↗ |

## .spec.template.spec.volumes[].persistentVolumeClaim

**Description**

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

**Type**

`object`

**Required**

`claimName`

| Property | Type | Description |
|---|---|---|
| `claimName` | `string` | claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims ↗ |
| `readOnly` | `boolean` | readOnly Will force the ReadOnly setting in VolumeMounts. Default false. |

## .spec.template.spec.volumes[].photonPersistentDisk

**Description**

Represents a Photon Controller persistent disk resource.

**Type**

`object`

**Required**

`pdID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `pdID` | `string` | pdID is the ID that identifies Photon Controller persistent disk |

## .spec.template.spec.volumes[].portworxVolume

**Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `volumeID` | `string` | volumeID uniquely identifies a Portworx volume |

## .spec.template.spec.volumes[].projected

**Description**

Represents a projected volume source

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `defaultMode` | `integer` | defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `sources` | `array` | sources is the list of volume projections. Each entry in this list handles one source. |

## .spec.template.spec.volumes[].projected.sources

**Description**

sources is the list of volume projections. Each entry in this list handles one source.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[]

**Description**

Projection that may be projected along with other supported volume types. Exactly one of these fields must be set.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `clusterTrustBundle` | `object` | ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem. |
| `configMap` | `object` | Adapts a ConfigMap into a projected volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode. |
| `downwardAPI` | `object` | Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode. |
| `secret` | `object` | Adapts a secret into a projected volume.<br><br>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode. |
| `serviceAccountToken` | `object` | ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise). |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle

**Description**

ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem.

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `name` | `string` | Select a single ClusterTrustBundle by object name. Mutually-exclusive with signerName and labelSelector. |
| `optional` | `boolean` | If true, don't block pod startup if the referenced ClusterTrustBundle(s) aren't available. If using name, then the named ClusterTrustBundle is allowed not to exist. If using signerName, then the combination of signerName and labelSelector is allowed to match zero ClusterTrustBundles. |
| `path` | `string` | Relative path from the volume root to write the bundle. |
| `signerName` | `string` | Select all ClusterTrustBundles that match this signer name. Mutually-exclusive with name. The contents of all selected ClusterTrustBundles will be unified and deduplicated. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].projected.sources[].configMap

**Description**

Adapts a ConfigMap into a projected volume. The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].projected.sources[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI

**Description**

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | Items is a list of DownwardAPIVolume file |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items

**Description**

Items is a list of DownwardAPIVolume file

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |

| Property | Type | Description |
| --- | --- | --- |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and |

| Property | Type | Description |
| --- | --- | --- |
| | | The serialization format is:<br>`(Note that <suffix> may be empty, from the "" case in <decimalSI>.)`<br>`<digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits`<br>`(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht`<br>`<decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E`<br>`(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)`<br>`<decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber>` ```<br>`No matter which of the three exponent forms is used, no quantity may represent a num`<br>`When a Quantity is parsed from a string, it will remember the type of suffix it had,`<br>`Before serializing, Quantity will be put in "canonical form". This means that Expone`<br>`- No precision is lost - No fractional digits will be emitted - The exponent (or suf`<br>`The sign will be omitted unless the number is negative.`<br>`Examples:`<br>`- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"`<br>`Note that the quantity will NEVER be internally represented by a floating point numb`<br>`Non-canonical values will still parse as long as they are well formed, but will be r`<br>`This format is intended to make it difficult to use these numbers without writing so` |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].projected.sources[].secret

### Description

Adapts a secret into a projected volume. The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional field specify whether the Secret or its key must be defined |

## .spec.template.spec.volumes[].projected.sources[].secret.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].serviceAccountToken

**Description**

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `audience` | `string` | audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver. |
| `expirationSeconds` | `integer` | expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours.Defaults to 1 hour and must be at least 10 minutes. |
| `path` | `string` | path is the path relative to the mount point of the file to project the token into. |

## .spec.template.spec.volumes[].quobyte

**Description**

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`registry` `volume`

| Property | Type | Description |
|---|---|---|
| `group` | `string` | group to map volume access to Default is no group |
| `readOnly` | `boolean` | readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false. |
| `registry` | `string` | registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes |
| `tenant` | `string` | tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin |
| `user` | `string` | user to map volume access to Defaults to serivceaccount user |

| Property | Type | Description |
| --- | --- | --- |
| `volume` | `string` | volume is a string that references an already created Quobyte volume by name. |

## .spec.template.spec.volumes[].rbd

### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`monitors` `image`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd ↗ |
| `image` | `string` | image is the rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `keyring` | `string` | keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `monitors` | `array` | monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `pool` | `string` | pool is the rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is the rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].rbd.monitors

**Description**

monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].rbd.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].rbd.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].scaleIO

**Description**

ScaleIOVolumeSource represents a persistent ScaleIO volume

**Type**

`object`

**Required**

`gateway` `system` `secretRef`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs". |
| `gateway` | `string` | gateway is the host address of the ScaleIO API Gateway. |
| `protectionDomain` | `string` | protectionDomain is the name of the ScaleIO Protection Domain for the configured storage. |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

| Property | Type | Description |
| --- | --- | --- |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `sslEnabled` | `boolean` | sslEnabled Flag enable/disable SSL communication with Gateway, default false |
| `storageMode` | `string` | storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned. |
| `storagePool` | `string` | storagePool is the ScaleIO Storage Pool associated with the protection domain. |
| `system` | `string` | system is the name of the storage system as configured in ScaleIO. |
| `volumeName` | `string` | volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source. |

## .spec.template.spec.volumes[].scaleIO.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].secret

**Description**

Adapts a Secret into a volume. The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON |

| Property | Type | Description |
| --- | --- | --- |
| | | requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `optional` | `boolean` | optional field specify whether the Secret or its keys must be defined |
| `secretName` | `string` | secretName is the name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret ↗ |

## .spec.template.spec.volumes[].secret.items

**Description**

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].storageos

### Description

Represents a StorageOS persistent volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeName` | `string` | volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace. |
| `volumeNamespace` | `string` | volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created. |

## .spec.template.spec.volumes[].storageos.secretRef

### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].vsphereVolume

**Description**

Represents a vSphere volume resource.

**Type**

`object`

**Required**

`volumePath`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `storagePolicyID` | `string` | storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName. |
| `storagePolicyName` | `string` | storagePolicyName is the storage Policy Based Management (SPBM) profile name. |
| `volumePath` | `string` | volumePath is the path that identifies vSphere volume vmdk |

## .status

**Description**

ReplicationControllerStatus represents the current status of a replication controller.

**Type**

`object`

**Required**

`replicas`

| Property | Type | Description |
|---|---|---|
| `availableReplicas` | `integer` | The number of available replicas (ready for at least minReadySeconds) for this replication controller. |
| `conditions` | `array` | Represents the latest available observations of a replication controller's current state. |
| `fullyLabeledReplicas` | `integer` | The number of pods that have labels matching the labels of the pod template of the replication controller. |
| `observedGeneration` | `integer` | ObservedGeneration reflects the generation of the most recently observed replication controller. |

| Property | Type | Description |
| --- | --- | --- |
| `readyReplicas` | `integer` | The number of ready replicas for this replication controller. |
| `replicas` | `integer` | Replicas is the most recently observed number of replicas. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#what-is-a-replicationcontroller ↗ |

## .status.conditions

**Description**

Represents the latest available observations of a replication controller's current state.

**Type**

`array`

## .status.conditions[]

**Description**

ReplicationControllerCondition describes the state of a replication controller at a certain point.

**Type**

`object`

**Required**

`type` `status`

| Property | Type | Description |
| --- | --- | --- |
| `lastTransitionTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | A human readable message indicating details about the transition. |
| `reason` | `string` | The reason for the condition's last transition. |
| `status` | `string` | Status of the condition, one of True, False, Unknown. |
| `type` | `string` | Type of replication controller condition. |

# API Endpoints

The following API endpoints are available:

- `/kubernetes/{cluster}/api/v1/namespaces/{namespace}/replicationcontrollers`
  - `DELETE` : delete collection of ReplicationController
  - `GET` : list objects of kind ReplicationController
  - `POST` : create a new ReplicationController
- `/kubernetes/{cluster}/api/v1/namespaces/{namespace}/replicationcontrollers/{name}`
  - `DELETE` : delete the specified ReplicationController
  - `GET` : read the specified ReplicationController
  - `PATCH` : partially update the specified ReplicationController
  - `PUT` : replace the specified ReplicationController
- `/kubernetes/{cluster}/api/v1/namespaces/{namespace}/replicationcontrollers/{name}/status`
  - `GET` : read status of the specified ReplicationController
  - `PATCH` : partially update status of the specified ReplicationController
  - `PUT` : replace status of the specified ReplicationController

## /kubernetes/{cluster}/api/v1/namespaces/{namespace}/replicationcontrollers

**HTTP method**

`DELETE`

**Description**

delete collection of ReplicationController

**HTTP responses**

| HTTP code | Response body |
|---|---|
| 200 - OK | `Status` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`GET`

**Description**

list objects of kind ReplicationController

**HTTP responses**

| HTTP code | Response body |
|---|---|
| 200 - OK | `ReplicationControllerList` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`POST`

**Description**

create a new ReplicationController

**Query parameters**

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
|---|---|---|
| `body` | `ReplicationController` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `ReplicationController` schema |
| 201 - Created | `ReplicationController` schema |
| 202 - Accepted | `ReplicationController` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/api/v1/namespaces/{namespace}/replicationcontrollers/{name}

### HTTP method

`DELETE`

### Description

delete the specified ReplicationController

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `Status` schema |
| 202 - Accepted | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

**Description**

read the specified ReplicationController

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicationController` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`PATCH`

**Description**

partially update the specified ReplicationController

**Query parameters**

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicationController` schema |
| 401 - Unauthorized | Empty |

**HTTP method**

`PUT`

**Description**

replace the specified ReplicationController

**Query parameters**

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default |

| Parameter | Type | Description |
|---|---|---|
| | | behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
|---|---|---|
| `body` | `ReplicationController` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `ReplicationController` schema |
| 201 - Created | `ReplicationController` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/api/v1/namespaces/{namespace}/replicationcontrollers/{name}/status

### HTTP method

`GET`

### Description

read status of the specified ReplicationController

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `ReplicationController` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update status of the specified ReplicationController

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each |

| Parameter | Type | Description |
| --- | --- | --- |
| | | unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicationController` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace status of the specified ReplicationController

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `ReplicationController` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `ReplicationController` schema |
| 201 - Created | `ReplicationController` schema |
| 401 - Unauthorized | Empty |

# StatefulSet [apps/v1]

**Description**

StatefulSet represents a set of pods with consistent identities. Identities are defined as: - Network: A single stable DNS and hostname. - Storage: As many VolumeClaims as requested. The StatefulSet guarantees that a given network identity will always map to the same storage identity.

**Type**

`object`

## Specification

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources |
| `kind` | `string` | Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | A StatefulSetSpec is the specification of a StatefulSet. |
| `status` | `object` | StatefulSetStatus represents the current state of a StatefulSet. |

### .spec

**Description**

A StatefulSetSpec is the specification of a StatefulSet.

**Type**

`object`

**Required**

`selector` `template` `serviceName`

| Property | Type | Description |
| --- | --- | --- |
| `minReadySeconds` | `integer` | Minimum number of seconds for which a newly created pod should be ready without any of its container crashing for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready) |
| `ordinals` | `object` | StatefulSetOrdinals describes the policy used for replica ordinal assignment in this StatefulSet. |
| `persistentVolumeClaimRetentionPolicy` | `object` | StatefulSetPersistentVolumeClaimRetentionPolicy describes the policy used for PVCs created from the StatefulSet VolumeClaimTemplates. |
| `podManagementPolicy` | `string` | podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is `OrderedReady`, where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is `Parallel` which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.<br><br>Possible enum values:<br><br>• `"OrderedReady"` will create pods in strictly increasing order on scale up and strictly decreasing order on scale down, progressing only when the previous pod is ready or terminated. At most one pod will be changed at any time.<br>• `"Parallel"` will create and delete pods as soon as the stateful set replica count is changed, and will not wait for pods to be ready or complete termination. |
| `replicas` | `integer` | replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1. |
| `revisionHistoryLimit` | `integer` | revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |

| Property | Type | Description |
| --- | --- | --- |
| `serviceName` | `string` | serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller. |
| `template` | `object` | PodTemplateSpec describes the data a pod should have when created from a template |
| `updateStrategy` | `object` | StatefulSetUpdateStrategy indicates the strategy that the StatefulSet controller will use to perform updates. It includes any additional parameters necessary to perform the update for the indicated strategy. |
| `volumeClaimTemplates` | `array` | volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name. |

## .spec.ordinals

**Description**

StatefulSetOrdinals describes the policy used for replica ordinal assignment in this StatefulSet.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `start` | `integer` | start is the number representing the first replica's index. It may be used to number replicas from an alternate index (eg: 1-indexed) over the default 0-indexed names, or to orchestrate progressive movement of replicas from one StatefulSet to another. If set, replica indices will be in the range: [.spec.ordinals.start, .spec.ordinals.start + .spec.replicas). If unset, defaults to 0. Replica indices will be in the range: [0, .spec.replicas). |

## .spec.persistentVolumeClaimRetentionPolicy

**Description**

StatefulSetPersistentVolumeClaimRetentionPolicy describes the policy used for PVCs created from the StatefulSet VolumeClaimTemplates.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `whenDeleted` | `string` | WhenDeleted specifies what happens to PVCs created from StatefulSet VolumeClaimTemplates when the StatefulSet is deleted. The default policy of `Retain` causes PVCs to not be affected by StatefulSet deletion. The `Delete` policy |

| Property | Type | Description |
| --- | --- | --- |
| | | causes those PVCs to be deleted. |
| `whenScaled` | `string` | WhenScaled specifies what happens to PVCs created from StatefulSet VolumeClaimTemplates when the StatefulSet is scaled down. The default policy of `Retain` causes PVCs to not be affected by a scaledown. The `Delete` policy causes the associated PVCs for any excess pods above the replica count to be deleted. |

## .spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |

| Property | Type | Description |
| --- | --- | --- |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template

**Description**

PodTemplateSpec describes the data a pod should have when created from a template

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PodSpec is a description of a pod. |

## .spec.template.spec

**Description**

PodSpec is a description of a pod.

**Type**

`object`

**Required**

`containers`

| Property | Type | Description |
| --- | --- | --- |
| `activeDeadlineSeconds` | `integer` | Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer. |
| `affinity` | `object` | Affinity is a group of affinity scheduling rules. |
| `automountServiceAccountToken` | `boolean` | AutomountServiceAccountToken indicates whether a service account token should be automatically mounted. |
| `containers` | `array` | List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated. |
| `dnsConfig` | `object` | PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy. |
| `dnsPolicy` | `string` | Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.<br><br>Possible enum values:<br>• `"ClusterFirst"` indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>• `"ClusterFirstWithHostNet"` indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings.<br>• `"Default"` indicates that the pod should use the default (as determined by kubelet) DNS settings.<br>• `"None"` indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig. |
| `enableServiceLinks` | `boolean` | EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching traditional container linking syntax. Optional: Defaults to true. |
| `ephemeralContainers` | `array` | List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating |

| Property | Type | Description |
| --- | --- | --- |
| | | a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. |
| `hostAliases` | `array` | HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. |
| `hostIPC` | `boolean` | Use the host's ipc namespace. Optional: Default to false. |
| `hostNetwork` | `boolean` | Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false. |
| `hostPID` | `boolean` | Use the host's pid namespace. Optional: Default to false. |
| `hostUsers` | `boolean` | Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new userns is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature. |
| `hostname` | `string` | Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value. |
| `imagePullSecrets` | `array` | ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod ↗ |
| `initContainers` | `array` | List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/ ↗ |
| `nodeName` | `string` | NodeName indicates in which node this pod is scheduled. If empty, this pod is a candidate for scheduling by the scheduler defined in schedulerName. Once this field is set, the kubelet for this node becomes responsible for the lifecycle of this pod. This field should not be used to express a |

| Property | Type | Description |
| --- | --- | --- |
| | | desire for the pod to be scheduled on a specific node. https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodename ↗ |
| `nodeSelector` | `object` | NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/ ↗ |
| `os` | `object` | PodOS defines the OS parameters of a pod. |
| `overhead` | `object` | Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md ↗ |
| `preemptionPolicy` | `string` | PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.<br><br>Possible enum values:<br>• `"Never"` means that pod never preempts other pods with lower priority.<br>• `"PreemptLowerPriority"` means that pod can preempt other pods with lower priority. |
| `priority` | `integer` | The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority. |
| `priorityClassName` | `string` | If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default. |
| `readinessGates` | `array` | If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates ↗ |
| `resourceClaims` | `array` | ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name. |

| Property | Type | Description |
| --- | --- | --- |
| | | This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br><br>This field is immutable. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy<br><br>Possible enum values:<br>- `"Always"`<br>- `"Never"`<br>- `"OnFailure"` |
| `runtimeClassName` | `string` | RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class |
| `schedulerName` | `string` | If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler. |
| `schedulingGates` | `array` | SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.<br><br>SchedulingGates can only be set at pod creation time, and be removed only afterwards. |
| `securityContext` | `object` | PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext. |
| `serviceAccount` | `string` | DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead. |
| `serviceAccountName` | `string` | ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/ |

| Property | Type | Description |
| --- | --- | --- |
| `setHostnameAsFQDN` | `boolean` | If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false. |
| `shareProcessNamespace` | `boolean` | Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false. |
| `subdomain` | `string` | If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all. |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds. |
| `tolerations` | `array` | If specified, the pod's tolerations. |
| `topologySpreadConstraints` | `array` | TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed. |
| `volumes` | `array` | List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes ↗ |

## .spec.template.spec.affinity

### Description

Affinity is a group of affinity scheduling rules.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `nodeAffinity` | `object` | Node affinity is a group of node affinity scheduling rules. |
| `podAffinity` | `object` | Pod affinity is a group of inter pod affinity scheduling rules. |
| `podAntiAffinity` | `object` | Pod anti affinity is a group of inter pod anti affinity scheduling rules. |

## .spec.template.spec.affinity.nodeAffinity

**Description**

Node affinity is a group of node affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `object` | A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

**Type**

`object`

**Required**

`weight` `preference`

| Property | Type | Description |
|---|---|---|
| `preference` | `object` | A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm. |
| `weight` | `integer` | Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

**Type**

`object`

**Required**

`nodeSelectorTerms`

| Property | Type | Description |
| --- | --- | --- |
| `nodeSelectorTerms` | `array` | Required. A list of node selector terms. The terms are ORed. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms

**Description**

Required. A list of node selector terms. The terms are ORed.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]

**Description**

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | A list of node selector requirements by node's labels. |
| `matchFields` | `array` | A list of node selector requirements by node's fields. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions

**Description**

A list of node selector requirements by node's labels.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields

**Description**

A list of node selector requirements by node's fields.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

**Description**

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The label key that the selector applies to. |
| `operator` | `string` | Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.<br><br>Possible enum values:<br>- `"DoesNotExist"`<br>- `"Exists"`<br>- `"Gt"`<br>- `"In"`<br>- `"Lt"`<br>- `"NotIn"` |
| `values` | `array` | An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

**Description**

An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity

**Description**

Pod affinity is a group of inter pod affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
|---|---|---|
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` |

| Property | Type | Description |
| --- | --- | --- |
| | | as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or

may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |

| Property | Type | Description |
| --- | --- | --- |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity

**Description**

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `preferredDuringSchedulingIgnoredDuringExecution` | `array` | The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred. |
| `requiredDuringSchedulingIgnoredDuringExecution` | `array` | If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution

**Description**

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

**Description**

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

**Type**

`object`

**Required**

`weight` `podAffinityTerm`

| Property | Type | Description |
|---|---|---|
| `podAffinityTerm` | `object` | Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key matches that of any node on which a pod of the set of pods is running |
| `weight` | `integer` | weight associated with matching the corresponding podAffinityTerm, in the range 1-100. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |

| Property | Type | Description |
| --- | --- | --- |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution

**Description**

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

**Description**

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

**Type**

`object`

**Required**

`topologyKey`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |
| `mismatchLabelKeys` | `array` | MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `namespaceSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `namespaces` | `array` | namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace". |
| `topologyKey` | `string` | This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key in (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both matchLabelKeys and labelSelector. Also, matchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

**Description**

MismatchLabelKeys is a set of pod label keys to select which pods will be taken into consideration. The keys are used to lookup values from the incoming pod labels, those key-value labels are merged with `labelSelector` as `key notin (value)` to select the group of existing pods which pods will be taken into consideration for the incoming pod's pod (anti) affinity. Keys that don't exist in the incoming pod labels will be ignored. The default value is empty. The same key is forbidden to exist in both mismatchLabelKeys and labelSelector. Also, mismatchLabelKeys cannot be set when labelSelector isn't set. This is a beta field and requires enabling MatchLabelKeysInPodAffinity feature gate (enabled by default).

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only |

| Property | Type | Description |
| --- | --- | --- |
| | | "value". The requirements are ANDed. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

**Description**

namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".

**Type**

`array`

## .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

**Type**

`string`

## .spec.template.spec.containers

**Description**

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[]

**Description**

A single application container that you want to run within a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |

| Property | Type | Description |
| --- | --- | --- |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |

| Property | Type | Description |
| --- | --- | --- |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br><br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.containers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].args[]

**Type**

`string`

## .spec.template.spec.containers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.containers[].command[]

**Type**

`string`

## .spec.template.spec.containers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.containers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.containers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.containers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.containers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.containers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.containers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.containers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.containers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.containers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.containers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.containers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.containers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].livenessProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.containers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.containers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.containers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.containers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.containers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.containers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.containers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.containers[].securityContext.appArmorProfile

### Description

AppArmorProfile defines a pod or container's AppArmor settings.

### Type

`object`

### Required

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.containers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.containers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.containers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.containers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.containers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.containers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa ↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.containers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.containers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.containers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.containers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.containers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.containers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.containers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.containers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.containers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.containers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.containers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.dnsConfig

**Description**

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `nameservers` | `array` | A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed. |
| `options` | `array` | A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy. |
| `searches` | `array` | A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed. |

## .spec.template.spec.dnsConfig.nameservers

**Description**

A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.nameservers[]

**Type**

`string`

## .spec.template.spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

`array`

## .spec.template.spec.dnsConfig.options[]

**Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is this DNS resolver option's name. Required. |
| `value` | `string` | Value is this DNS resolver option's value. |

## .spec.template.spec.dnsConfig.searches

**Description**

A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

**Type**

`array`

## .spec.template.spec.dnsConfig.searches[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers

**Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[]

**Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable |

| Property | Type | Description |
| --- | --- | --- |
| | | exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `command` | `array` | Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | Ports are not allowed for ephemeral containers. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | Restart policy for the container to manage the restart behavior of each container within a pod. This may only be set for init containers. You cannot set this field on ephemeral containers. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |
| `targetContainerName` | `string` | If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.<br><br>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined. |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message |

| Property | Type | Description |
| --- | --- | --- |
| | | length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>- `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>- `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.ephemeralContainers[].args

**Description**

Arguments to the entrypoint. The image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].args[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.ephemeralContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.ephemeralContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.ephemeralContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].ports

**Description**

Ports are not allowed for ephemeral containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |
| `hostIP` | `string` | What host IP to bind the external port to. |

| Property | Type | Description |
| --- | --- | --- |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br><br>Possible enum values:<br>- `"SCTP"` is the SCTP protocol.<br>- `"TCP"` is the TCP protocol.<br>- `"UDP"` is the UDP protocol. |

## .spec.template.spec.ephemeralContainers[].readinessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.ephemeralContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.ephemeralContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.ephemeralContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.ephemeralContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>- `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>- `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.ephemeralContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.ephemeralContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.ephemeralContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.ephemeralContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.ephemeralContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.ephemeralContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.ephemeralContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.ephemeralContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

**Type**

`array`

## .spec.template.spec.ephemeralContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.hostAliases

**Description**

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified.

**Type**

`array`

## .spec.template.spec.hostAliases[]

**Description**

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

**Type**

`object`

**Required**

`ip`

| Property | Type | Description |
|---|---|---|
| `hostnames` | `array` | Hostnames for the above IP address. |
| `ip` | `string` | IP address of the host file entry. |

## .spec.template.spec.hostAliases[].hostnames

**Description**

Hostnames for the above IP address.

**Type**

`array`

## .spec.template.spec.hostAliases[].hostnames[]

**Type**

`string`

## .spec.template.spec.imagePullSecrets

**Description**

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

**Type**

`array`

## .spec.template.spec.imagePullSecrets[]

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.initContainers

### Description

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

### Type

`array`

## .spec.template.spec.initContainers[]

### Description

A single application container that you want to run within a pod.

### Type

`object`

### Required

`name`

| Property | Type | Description |
| --- | --- | --- |
| `args` | `array` | Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell ↗ |
| `env` | `array` | List of environment variables to set in the container. Cannot be updated. |
| `envFrom` | `array` | List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated. |
| `image` | `string` | Container image name. More info: https://kubernetes.io/docs/concepts/containers/images ↗ This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |
| `imagePullPolicy` | `string` | Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images ↗<br><br>Possible enum values:<br>• `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>• `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>• `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `lifecycle` | `object` | Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted. |
| `livenessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `name` | `string` | Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. |

| Property | Type | Description |
| --- | --- | --- |
| `ports` | `array` | List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255 ↗. Cannot be updated. |
| `readinessProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `resizePolicy` | `array` | Resources resize policy for the container. |
| `resources` | `object` | ResourceRequirements describes the compute resource requirements. |
| `restartPolicy` | `string` | RestartPolicy defines the restart behavior of individual containers in a pod. This field may only be set for init containers, and the only allowed value is "Always". For non-init containers or when this field is not specified, the restart behavior is defined by the Pod's restart policy and the container type. Setting the RestartPolicy as "Always" for the init container will have the following effect: this init container will be continually restarted on exit until all regular containers have terminated. Once all regular containers have completed, all init containers with restartPolicy "Always" will be shut down. This lifecycle differs from normal init containers and is often referred to as a "sidecar" container. Although this init container still starts in the init container sequence, it does not wait for the container to complete before proceeding to the next init container. Instead, the next init container starts immediately after this init container is started, or after any startupProbe has successfully completed. |
| `securityContext` | `object` | SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence. |
| `startupProbe` | `object` | Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic. |
| `stdin` | `boolean` | Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false. |
| `stdinOnce` | `boolean` | Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false |

| Property | Type | Description |
| --- | --- | --- |
| `terminationMessagePath` | `string` | Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated. |
| `terminationMessagePolicy` | `string` | Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.<br>Possible enum values:<br>• `"FallbackToLogsOnError"` will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.<br>• `"File"` is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. |
| `tty` | `boolean` | Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false. |
| `volumeDevices` | `array` | volumeDevices is the list of block devices to be used by the container. |
| `volumeMounts` | `array` | Pod volumes to mount into the container's filesystem. Cannot be updated. |
| `workingDir` | `string` | Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated. |

## .spec.template.spec.initContainers[].args

**Description**

Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].args[]

**Type**

`string`

## .spec.template.spec.initContainers[].command

**Description**

Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

**Type**

`array`

## .spec.template.spec.initContainers[].command[]

**Type**

`string`

## .spec.template.spec.initContainers[].env

**Description**

List of environment variables to set in the container. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].env[]

**Description**

EnvVar represents an environment variable present in a Container.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name of the environment variable. Must be a C_IDENTIFIER. |
| `value` | `string` | Variable references $(VAR_NAME) are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double $$ are reduced to a single $, which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will produce the string literal "$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "". |
| `valueFrom` | `object` | EnvVarSource represents a source for the value of an EnvVar. |

## .spec.template.spec.initContainers[].env[].valueFrom

**Description**

EnvVarSource represents a source for the value of an EnvVar.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `configMapKeyRef` | `object` | Selects a key from a ConfigMap. |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |
| `secretKeyRef` | `object` | SecretKeySelector selects a key of a Secret. |

## .spec.template.spec.initContainers[].env[].valueFrom.configMapKeyRef

**Description**

Selects a key from a ConfigMap.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | The key to select. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | Specify whether the ConfigMap or its key must be defined |

## .spec.template.spec.initContainers[].env[].valueFrom.fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.initContainers[].env[].valueFrom.resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.initContainers[].env[].valueFrom.secretKeyRef

**Description**

SecretKeySelector selects a key of a Secret.

**Type**

`object`

**Required**

`key`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | The key of the secret to select from. Must be a valid secret key. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret or its key must be defined |

## .spec.template.spec.initContainers[].envFrom

**Description**

List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `configMapRef` | `object` | ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.<br><br>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables. |
| `prefix` | `string` | An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER. |
| `secretRef` | `object` | SecretEnvSource selects a Secret to populate the environment variables with.<br><br>The contents of the target Secret's Data field will represent the key-value pairs as environment variables. |

## .spec.template.spec.initContainers[].envFrom[].configMapRef

**Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the ConfigMap must be defined |

## .spec.template.spec.initContainers[].envFrom[].secretRef

**Description**

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | Specify whether the Secret must be defined |

## .spec.template.spec.initContainers[].lifecycle

**Description**

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `postStart` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |
| `preStop` | `object` | LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified. |

## .spec.template.spec.initContainers[].lifecycle.postStart

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.postStart.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].lifecycle.postStart.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.postStart.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.postStart.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].lifecycle.preStop

**Description**

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `sleep` | `object` | SleepAction describes a "sleep" action. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].lifecycle.preStop.sleep

**Description**

SleepAction describes a "sleep" action.

**Type**

`object`

**Required**

`seconds`

| Property | Type | Description |
| --- | --- | --- |
| `seconds` | `integer` | Seconds is the number of seconds to sleep. |

## .spec.template.spec.initContainers[].lifecycle.preStop.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].livenessProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].livenessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].livenessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].livenessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].livenessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br><br>Possible enum values:<br>• `"HTTP"` means that the scheme used will be http://<br>• `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |

| Property | Type | Description |
| --- | --- | --- |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].livenessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].ports

**Description**

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See https://github.com/kubernetes/kubernetes/issues/108255. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

`object`

**Required**

`containerPort`

| Property | Type | Description |
| --- | --- | --- |
| `containerPort` | `integer` | Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536. |

| Property | Type | Description |
| --- | --- | --- |
| `hostIP` | `string` | What host IP to bind the external port to. |
| `hostPort` | `integer` | Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. |
| `name` | `string` | If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. |
| `protocol` | `string` | Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".<br>Possible enum values:<br>• `"SCTP"` is the SCTP protocol.<br>• `"TCP"` is the TCP protocol.<br>• `"UDP"` is the UDP protocol. |

## .spec.template.spec.initContainers[].readinessProbe

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |

| Property | Type | Description |
| --- | --- | --- |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].readinessProbe.exec

**Description**

ExecAction describes a "run in container" action.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].readinessProbe.exec.command

**Description**

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].readinessProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].readinessProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].readinessProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].resizePolicy

**Description**

Resources resize policy for the container.

**Type**

`array`

## .spec.template.spec.initContainers[].resizePolicy[]

**Description**

ContainerResizePolicy represents resource resize policy for the container.

**Type**

`object`

**Required**

`resourceName` `restartPolicy`

| Property | Type | Description |
| --- | --- | --- |
| `resourceName` | `string` | Name of the resource to which this resource resize policy applies. Supported values: cpu, memory. |
| `restartPolicy` | `string` | Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired. |

## .spec.template.spec.initContainers[].resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ |

## .spec.template.spec.initContainers[].resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.initContainers[].resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.initContainers[].resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.initContainers[].securityContext

**Description**

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `allowPrivilegeEscalation` | `boolean` | AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows. |
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `capabilities` | `object` | Adds and removes POSIX capabilities from running containers. |
| `privileged` | `boolean` | Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows. |
| `procMount` | `string` | procMount denotes the type of proc mount to use for the containers. The default value is Default which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br>• `"Default"` uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information.<br>• `"Unmasked"` bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications. |
| `readOnlyRootFilesystem` | `boolean` | Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.initContainers[].securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br><br>- `"Localhost"` indicates that a profile pre-loaded on the node should be used. |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.initContainers[].securityContext.capabilities

**Description**

Adds and removes POSIX capabilities from running containers.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `add` | `array` | Added capabilities |
| `drop` | `array` | Removed capabilities |

## .spec.template.spec.initContainers[].securityContext.capabilities.add

**Description**

Added capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.add[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop

**Description**

Removed capabilities

**Type**

`array`

## .spec.template.spec.initContainers[].securityContext.capabilities.drop[]

**Type**

`string`

## .spec.template.spec.initContainers[].securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `level` | `string` | Level is SELinux level label that applies to the container. |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.initContainers[].securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
|---|---|---|
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.initContainers[].securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.initContainers[].startupProbe

**Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `exec` | `object` | ExecAction describes a "run in container" action. |
| `failureThreshold` | `integer` | Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. |
| `grpc` | `object` | GRPCAction specifies an action involving a GRPC service. |
| `httpGet` | `object` | HTTPGetAction describes an action based on HTTP Get requests. |
| `initialDelaySeconds` | `integer` | Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes |

| Property | Type | Description |
| --- | --- | --- |
| `periodSeconds` | `integer` | How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. |
| `successThreshold` | `integer` | Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1. |
| `tcpSocket` | `object` | TCPSocketAction describes an action based on opening a socket |
| `terminationGracePeriodSeconds` | `integer` | Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset. |
| `timeoutSeconds` | `integer` | Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes ↗ |

## .spec.template.spec.initContainers[].startupProbe.exec

### Description

ExecAction describes a "run in container" action.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `command` | `array` | Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('\|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy. |

## .spec.template.spec.initContainers[].startupProbe.exec.command

### Description

Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions ('|', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

### Type

`array`

## .spec.template.spec.initContainers[].startupProbe.exec.command[]

**Type**

`string`

## .spec.template.spec.initContainers[].startupProbe.grpc

**Description**

GRPCAction specifies an action involving a GRPC service.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `port` | `integer` | Port number of the gRPC service. Number must be in the range 1 to 65535. |
| `service` | `string` | Service is the name of the service to place in the gRPC HealthCheckRequest (see https://github.com/grpc/grpc/blob/master/doc/health-checking.md ↗).<br><br>If this is not specified, the default behavior is defined by gRPC. |

## .spec.template.spec.initContainers[].startupProbe.httpGet

**Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
|---|---|---|
| `host` | `string` | Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead. |
| `httpHeaders` | `array` | Custom headers to set in the request. HTTP allows repeated headers. |
| `path` | `string` | Path to access on the HTTP server. |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

| Property | Type | Description |
| --- | --- | --- |
| `scheme` | `string` | Scheme to use for connecting to the host. Defaults to HTTP.<br>Possible enum values:<br>- `"HTTP"` means that the scheme used will be http://<br>- `"HTTPS"` means that the scheme used will be https:// |

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders

**Description**

Custom headers to set in the request. HTTP allows repeated headers.

**Type**

`array`

## .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]

**Description**

HTTPHeader describes a custom header to be used in HTTP probes

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header. |
| `value` | `string` | The header field value |

## .spec.template.spec.initContainers[].startupProbe.tcpSocket

**Description**

TCPSocketAction describes an action based on opening a socket

**Type**

`object`

**Required**

`port`

| Property | Type | Description |
| --- | --- | --- |
| `host` | `string` | Optional: Host name to connect to, defaults to the pod IP. |

| Property | Type | Description |
| --- | --- | --- |
| `port` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |

## .spec.template.spec.initContainers[].volumeDevices

**Description**

volumeDevices is the list of block devices to be used by the container.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeDevices[]

**Description**

volumeDevice describes a mapping of a raw block device within a container.

**Type**

`object`

**Required**

`name` `devicePath`

| Property | Type | Description |
| --- | --- | --- |
| `devicePath` | `string` | devicePath is the path inside of the container that the device will be mapped to. |
| `name` | `string` | name must match the name of a persistentVolumeClaim in the pod |

## .spec.template.spec.initContainers[].volumeMounts

**Description**

Pod volumes to mount into the container's filesystem. Cannot be updated.

**Type**

`array`

## .spec.template.spec.initContainers[].volumeMounts[]

**Description**

VolumeMount describes a mounting of a Volume within a container.

**Type**

`object`

**Required**

`name` `mountPath`

| Property | Type | Description |
| --- | --- | --- |
| `mountPath` | `string` | Path within the container at which the volume should be mounted. Must not contain ':'. |
| `mountPropagation` | `string` | mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10. When RecursiveReadOnly is set to IfPossible or to Enabled, MountPropagation must be None or unspecified (which defaults to None).<br><br>Possible enum values:<br><br>• `"Bidirectional"` means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology).<br>• `"HostToContainer"` means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).<br>• `"None"` means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology. |
| `name` | `string` | This must match the Name of a Volume. |
| `readOnly` | `boolean` | Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false. |
| `recursiveReadOnly` | `string` | RecursiveReadOnly specifies whether read-only mounts should be handled recursively.<br><br>If ReadOnly is false, this field has no meaning and must be unspecified.<br><br>If ReadOnly is true, and this field is set to Disabled, the mount is not made recursively read-only. If this field is set to IfPossible, the mount is made recursively read-only, if it is supported by the container runtime. If this field is set to Enabled, the mount is made recursively read-only if it is supported by the container runtime, otherwise the pod will not be started and an error will be generated to indicate the reason.<br><br>If this field is set to IfPossible or Enabled, MountPropagation must be set to None (or be unspecified, which defaults to None).<br><br>If this field is not specified, it is treated as an equivalent of Disabled. |
| `subPath` | `string` | Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root). |
| `subPathExpr` | `string` | Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references $(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive. |

## .spec.template.spec.nodeSelector

**Description**

NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

**Type**

`object`

## .spec.template.spec.os

**Description**

PodOS defines the OS parameters of a pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration ↗ Clients should expect to handle additional values and treat unrecognized values in this field as os: null |

## .spec.template.spec.overhead

**Description**

Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md

**Type**

`object`

## .spec.template.spec.readinessGates

**Description**

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates

**Type**

`array`

## .spec.template.spec.readinessGates[]

**Description**

PodReadinessGate contains the reference to a pod condition

**Type**

`object`

**Required**

`conditionType`

| Property | Type | Description |
| --- | --- | --- |
| `conditionType` | `string` | ConditionType refers to a condition in the pod's condition list with matching type. |

## .spec.template.spec.resourceClaims

**Description**

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable.

**Type**

`array`

## .spec.template.spec.resourceClaims[]

**Description**

PodResourceClaim references exactly one ResourceClaim, either directly or by naming a ResourceClaimTemplate which is then turned into a ResourceClaim for the pod. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL. |
| `resourceClaimName` | `string` | ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |
| `resourceClaimTemplateName` | `string` | ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.<br><br>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The pod name and resource name, along with a generated component, will be used to form a unique name for the ResourceClaim, which will be recorded in pod.status.resourceClaimStatuses.<br><br>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.<br><br>Exactly one of ResourceClaimName and ResourceClaimTemplateName must be set. |

## .spec.template.spec.resources

**Description**

ResourceRequirements describes the compute resource requirements.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `claims` | `array` | Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.<br>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.<br>This field is immutable. It can only be set for containers. |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.template.spec.resources.claims

**Description**

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. This field is immutable. It can only be set for containers.

**Type**

`array`

## .spec.template.spec.resources.claims[]

**Description**

ResourceClaim references one entry in PodSpec.ResourceClaims.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
|---|---|---|
| `name` | `string` | Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container. |

| Property | Type | Description |
| --- | --- | --- |
| `request` | `string` | Request is the name chosen for a request in the referenced claim. If empty, everything from the claim is made available, otherwise only the result of this request. |

## .spec.template.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.schedulingGates

**Description**

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod. SchedulingGates can only be set at pod creation time, and be removed only afterwards.

**Type**

`array`

## .spec.template.spec.schedulingGates[]

**Description**

PodSchedulingGate is associated to a Pod to guard its scheduling.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the scheduling gate. Each scheduling gate must have a unique name field. |

## .spec.template.spec.securityContext

**Description**

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `appArmorProfile` | `object` | AppArmorProfile defines a pod or container's AppArmor settings. |
| `fsGroup` | `integer` | A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:<br>1. The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw----<br>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows. |
| `fsGroupChangePolicy` | `string` | fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership(and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.<br>Possible enum values:<br>- `"Always"` indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior.<br>- `"OnRootMismatch"` indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume. |
| `runAsGroup` | `integer` | The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |
| `runAsNonRoot` | `boolean` | Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |
| `runAsUser` | `integer` | The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows. |

| Property | Type | Description |
| --- | --- | --- |
| `seLinuxChangePolicy` | `string` | seLinuxChangePolicy defines how the container's SELinux label is applied to all volumes used by the Pod. It has no effect on nodes that do not support SELinux or to volumes does not support SELinux. Valid values are "MountOption" and "Recursive".<br><br>"Recursive" means relabeling of all files on all Pod volumes by the container runtime. This may be slow for large volumes, but allows mixing privileged and unprivileged Pods sharing the same volume on the same node.<br><br>"MountOption" mounts all eligible Pod volumes with `-o context` mount option. This requires all Pods that share the same volume to use the same SELinux label. It is not possible to share the same volume among privileged and unprivileged Pods. Eligible volumes are in-tree FibreChannel and iSCSI volumes, and all CSI volumes whose CSI driver announces SELinux support by setting spec.seLinuxMount: true in their CSIDriver instance. Other volumes are always re-labelled recursively. "MountOption" value is allowed only when SELinuxMount feature gate is enabled.<br><br>If not specified and SELinuxMount feature gate is enabled, "MountOption" is used. If not specified and SELinuxMount feature gate is disabled, "MountOption" is used for ReadWriteOncePod volumes and "Recursive" for all other volumes.<br><br>This field affects only Pods that have SELinux label set, either in PodSecurityContext or in SecurityContext of all containers.<br><br>All Pods that use the same volume should use the same seLinuxChangePolicy, otherwise some pods can get stuck in ContainerCreating state. Note that this field cannot be set when spec.os.name is windows. |
| `seLinuxOptions` | `object` | SELinuxOptions are the labels to be applied to the container |
| `seccompProfile` | `object` | SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set. |
| `supplementalGroups` | `array` | A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows. |
| `supplementalGroupsPolicy` | `string` | Defines how supplemental groups of the first container processes are calculated. Valid values are "Merge" and "Strict". If not specified, "Merge" is used. (Alpha) Using the field requires the SupplementalGroupsPolicy feature gate to be enabled and the container runtime must implement support for this feature. Note that this field cannot be set when spec.os.name is windows.<br><br>Possible enum values:<br><br>- `"Merge"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be merged with the primary user's groups as defined in the container image (in /etc/group). |

| Property | Type | Description |
| --- | --- | --- |
| | | • `"Strict"` means that the container's provided SupplementalGroups and FsGroup (specified in SecurityContext) will be used instead of any groups defined in the container image. |
| `sysctls` | `array` | Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows. |
| `windowsOptions` | `object` | WindowsSecurityContextOptions contain Windows-specific options and credentials. |

## .spec.template.spec.securityContext.appArmorProfile

**Description**

AppArmorProfile defines a pod or container's AppArmor settings.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile loaded on the node that should be used. The profile must be preconfigured on the node to work. Must match the loaded name of the profile. Must be set if and only if type is "Localhost". |
| `type` | `string` | type indicates which kind of AppArmor profile will be applied. Valid options are: Localhost - a profile pre-loaded on the node. RuntimeDefault - the container runtime's default profile. Unconfined - no AppArmor enforcement.<br><br>Possible enum values:<br>• `"Localhost"` indicates that a profile pre-loaded on the node should be used.<br>• `"RuntimeDefault"` indicates that the container runtime's default AppArmor profile should be used.<br>• `"Unconfined"` indicates that no AppArmor profile should be enforced. |

## .spec.template.spec.securityContext.seLinuxOptions

**Description**

SELinuxOptions are the labels to be applied to the container

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `level` | `string` | Level is SELinux level label that applies to the container. |

| Property | Type | Description |
| --- | --- | --- |
| `role` | `string` | Role is a SELinux role label that applies to the container. |
| `type` | `string` | Type is a SELinux type label that applies to the container. |
| `user` | `string` | User is a SELinux user label that applies to the container. |

## .spec.template.spec.securityContext.seccompProfile

**Description**

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

**Type**

`object`

**Required**

`type`

| Property | Type | Description |
| --- | --- | --- |
| `localhostProfile` | `string` | localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must be set if type is "Localhost". Must NOT be set for any other type. |
| `type` | `string` | type indicates which kind of seccomp profile will be applied. Valid options are:<br><br>Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.<br><br>Possible enum values:<br>• `"Localhost"` indicates a profile defined in a file on the node should be used. The file's location relative to /seccomp.<br>• `"RuntimeDefault"` represents the default container runtime seccomp profile.<br>• `"Unconfined"` indicates no seccomp profile is applied (A.K.A. unconfined). |

## .spec.template.spec.securityContext.supplementalGroups

**Description**

A list of groups applied to the first process run in each container, in addition to the container's primary GID and fsGroup (if specified). If the SupplementalGroupsPolicy feature is enabled, the supplementalGroupsPolicy field determines whether these are in addition to or instead of any group memberships defined in the container image. If unspecified, no additional groups are added, though group memberships defined in the container image may still be used, depending on the supplementalGroupsPolicy field. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.supplementalGroups[]

**Type**

`integer`

## .spec.template.spec.securityContext.sysctls

**Description**

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

**Type**

`array`

## .spec.template.spec.securityContext.sysctls[]

**Description**

Sysctl defines a kernel parameter to be set

**Type**

`object`

**Required**

`name` `value`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of a property to set |
| `value` | `string` | Value of a property to set |

## .spec.template.spec.securityContext.windowsOptions

**Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `gmsaCredentialSpec` | `string` | GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa↗) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field. |
| `gmsaCredentialSpecName` | `string` | GMSACredentialSpecName is the name of the GMSA credential spec to use. |

| Property | Type | Description |
| --- | --- | --- |
| `hostProcess` | `boolean` | HostProcess determines if a container should be run as a 'Host Process' container. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true. |
| `runAsUserName` | `string` | The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. |

## .spec.template.spec.tolerations

**Description**

If specified, the pod's tolerations.

**Type**

`array`

## .spec.template.spec.tolerations[]

**Description**

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `effect` | `string` | Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.<br><br>Possible enum values:<br>- `"NoExecute"` Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController.<br>- `"NoSchedule"` Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler.<br>- `"PreferNoSchedule"` Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler. |
| `key` | `string` | Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys. |
| `operator` | `string` | Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category. |

| Property | Type | Description |
| --- | --- | --- |
| | | Possible enum values:<br>- `"Equal"`<br>- `"Exists"` |
| `tolerationSeconds` | `integer` | TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system. |
| `value` | `string` | Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string. |

## .spec.template.spec.topologySpreadConstraints

**Description**

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[]

**Description**

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

**Type**

`object`

**Required**

`maxSkew` `topologyKey` `whenUnsatisfiable`

| Property | Type | Description |
| --- | --- | --- |
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `matchLabelKeys` | `array` | MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.<br><br>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default). |

| Property | Type | Description |
| --- | --- | --- |
| `maxSkew` | `integer` | MaxSkew describes the degree to which pods may be unevenly distributed. When `whenUnsatisfiable=DoNotSchedule`, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1. \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P \| - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When `whenUnsatisfiable=ScheduleAnyway`, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed. |
| `minDomains` | `integer` | MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.<br><br>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2: \| zone1 \| zone2 \| zone3 \| \| P P \| P P \| P P \| The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew. |
| `nodeAffinityPolicy` | `string` | NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.<br><br>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br><br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |
| `nodeTaintsPolicy` | `string` | NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.<br><br>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.<br><br>Possible enum values:<br><br>• `"Honor"` means use this scheduling directive when calculating pod topology spread skew.<br>• `"Ignore"` means ignore this scheduling directive when calculating pod topology spread skew. |

| Property | Type | Description |
| --- | --- | --- |
| `topologyKey` | `string` | TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field. |
| `whenUnsatisfiable` | `string` | WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: \| zone1 \| zone2 \| zone3 \| \| P P P \| P \| P \| If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it *more* imbalanced. It's a required field.<br><br>Possible enum values:<br>- `"DoNotSchedule"` instructs the scheduler not to schedule the pod when constraints are not satisfied.<br>- `"ScheduleAnyway"` instructs the scheduler to schedule the pod even if constraints are not satisfied. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.topologySpreadConstraints[].labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys

**Description**

MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector. This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).

**Type**

`array`

## .spec.template.spec.topologySpreadConstraints[].matchLabelKeys[]

**Type**

`string`

## .spec.template.spec.volumes

**Description**

List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

**Type**

`array`

## .spec.template.spec.volumes[]

**Description**

Volume represents a named volume in a pod that may be accessed by any container in the pod.

**Type**

`object`

**Required**

`name`

| Property | Type | Description |
| --- | --- | --- |
| `awsElasticBlockStore` | `object` | Represents a Persistent Disk resource in AWS.<br><br>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling. |
| `azureDisk` | `object` | AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod. |
| `azureFile` | `object` | AzureFile represents an Azure File Service mount on the host and bind mount to the pod. |
| `cephfs` | `object` | Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling. |
| `cinder` | `object` | Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership |

| Property | Type | Description |
| --- | --- | --- |
| | | management and SELinux relabeling. |
| `configMap` | `object` | Adapts a ConfigMap into a volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling. |
| `csi` | `object` | Represents a source location of a volume to mount, managed by an external CSI driver |
| `downwardAPI` | `object` | DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling. |
| `emptyDir` | `object` | Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling. |
| `ephemeral` | `object` | Represents an ephemeral volume that is handled by a normal storage driver. |
| `fc` | `object` | Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling. |
| `flexVolume` | `object` | FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin. |
| `flocker` | `object` | Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling. |
| `gcePersistentDisk` | `object` | Represents a Persistent Disk resource in Google Compute Engine.<br><br>A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling. |
| `gitRepo` | `object` | Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.<br><br>DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container. |

| Property | Type | Description |
| --- | --- | --- |
| `glusterfs` | `object` | Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling. |
| `hostPath` | `object` | Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling. |
| `image` | `object` | ImageVolumeSource represents a image volume resource. |
| `iscsi` | `object` | Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling. |
| `name` | `string` | name of the volume. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `nfs` | `object` | Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling. |
| `persistentVolumeClaim` | `object` | PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system). |
| `photonPersistentDisk` | `object` | Represents a Photon Controller persistent disk resource. |
| `portworxVolume` | `object` | PortworxVolumeSource represents a Portworx volume resource. |
| `projected` | `object` | Represents a projected volume source |
| `quobyte` | `object` | Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling. |
| `rbd` | `object` | Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling. |
| `scaleIO` | `object` | ScaleIOVolumeSource represents a persistent ScaleIO volume |

| Property | Type | Description |
| --- | --- | --- |
| `secret` | `object` | Adapts a Secret into a volume.<br><br>The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling. |
| `storageos` | `object` | Represents a StorageOS persistent volume resource. |
| `vsphereVolume` | `object` | Represents a vSphere volume resource. |

## .spec.template.spec.volumes[].awsElasticBlockStore

**Description**

Represents a Persistent Disk resource in AWS. An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). |
| `readOnly` | `boolean` | readOnly value true will force the readOnly setting in VolumeMounts. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |
| `volumeID` | `string` | volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore ↗ |

## .spec.template.spec.volumes[].azureDisk

**Description**

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`diskName` `diskURI`

| Property | Type | Description |
|---|---|---|
| `cachingMode` | `string` | cachingMode is the Host Caching mode: None, Read Only, Read Write.<br>Possible enum values:<br>• `"None"`<br>• `"ReadOnly"`<br>• `"ReadWrite"` |
| `diskName` | `string` | diskName is the Name of the data disk in the blob storage |
| `diskURI` | `string` | diskURI is the URI of data disk in the blob storage |
| `fsType` | `string` | fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `kind` | `string` | kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared<br>Possible enum values:<br>• `"Dedicated"`<br>• `"Managed"`<br>• `"Shared"` |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

## .spec.template.spec.volumes[].azureFile

**Description**

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

**Type**

`object`

**Required**

`secretName` `shareName`

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretName` | `string` | secretName is the name of secret that contains Azure Storage Account Name and Key |
| `shareName` | `string` | shareName is the azure share Name |

## .spec.template.spec.volumes[].cephfs

**Description**

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`monitors`

| Property | Type | Description |
| --- | --- | --- |
| `monitors` | `array` | monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `path` | `string` | path is Optional: Used as the mounted root, rather than the full Ceph tree, default is / |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretFile` | `string` | secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].cephfs.monitors

**Description**

monitors is Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].cephfs.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].cephfs.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].cinder

**Description**

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeID` | `string` | volumeID used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md ↗ |

## .spec.template.spec.volumes[].cinder.secretRef

### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].configMap

### Description

Adapts a ConfigMap into a volume. The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].configMap.items

### Description

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].csi

**Description**

Represents a source location of a volume to mount, managed by an external CSI driver

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
|---|---|---|
| `driver` | `string` | driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster. |
| `fsType` | `string` | fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply. |

| Property | Type | Description |
| --- | --- | --- |
| `nodePublishSecretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `readOnly` | `boolean` | readOnly specifies a read-only configuration for the volume. Defaults to false (read/write). |
| `volumeAttributes` | `object` | volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values. |

## .spec.template.spec.volumes[].csi.nodePublishSecretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].csi.volumeAttributes

**Description**

volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**Type**

`object`

## .spec.template.spec.volumes[].downwardAPI

**Description**

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | Items is a list of downward API volume file |

## .spec.template.spec.volumes[].downwardAPI.items

**Description**

Items is a list of downward API volume file

**Type**

`array`

## .spec.template.spec.volumes[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
|---|---|---|
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
|---|---|---|
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and<br><br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].emptyDir

### Description

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `medium` | `string` | medium represents what type of storage medium should back this directory. The default is "" which means to use the n |
| `sizeLimit` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAM<br>The serialization format is: |

| Property | Type | Description |
| --- | --- | --- |
| | | (Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits> <r<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.html)<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a number<br>When a Quantity is parsed from a string, it will remember the type of suffix it had, an<br>Before serializing, Quantity will be put in "canonical form". This means that Exponent/<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suffix<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point number.<br>Non-canonical values will still parse as long as they are well formed, but will be re-e<br>This format is intended to make it difficult to use these numbers without writing some |

## .spec.template.spec.volumes[].ephemeral

**Description**

Represents an ephemeral volume that is handled by a normal storage driver.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `volumeClaimTemplate` | `object` | PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate

**Description**

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

**Type**

`object`

**Required**

`spec`

| Property | Type | Description |
| --- | --- | --- |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |
| `spec` | `object` | PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec

### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `accessModes` | `array` | accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1 ↗ |
| `dataSource` | `object` | TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace. |
| `dataSourceRef` | `object` | TypedObjectReference contains enough information to let you locate the typed referenced object |
| `resources` | `object` | VolumeResourceRequirements describes the storage resource requirements for a volume. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `storageClassName` | `string` | storageClassName is the name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1 ↗ |
| `volumeAttributesClassName` | `string` | volumeAttributesClassName may be used to set the VolumeAttributesClass used by this claim. If specified, the CSI driver will create or update the volume with the attributes defined in the corresponding VolumeAttributesClass. This has a different purpose than storageClassName, it can be changed after the claim is created. An empty string value means that no VolumeAttributesClass will be applied to the claim but it's not allowed to reset this field to empty string once it is set. If unspecified and the PersistentVolumeClaim is unbound, the default VolumeAttributesClass will be set by the persistentvolume controller if it exists. If the resource referred to by volumeAttributesClass does not exist, this PersistentVolumeClaim will be set to a Pending state, as reflected by the |

| Property | Type | Description |
| --- | --- | --- |
| | | modifyVolumeStatus field, until such as a resource exists. More info: https://kubernetes.io/docs/concepts/storage/volume-attributes-classes/ (Beta) Using this field requires the VolumeAttributesClass feature gate to be enabled (off by default). |
| `volumeMode` | `string` | volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.<br><br>Possible enum values:<br>• `"Block"` means the volume will not be formatted with a filesystem and will remain a raw block device.<br>• `"Filesystem"` means the volume will be or is formatted with a filesystem. |
| `volumeName` | `string` | volumeName is the binding reference to the PersistentVolume backing this claim. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes

**Description**

accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.accessModes[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

**Description**

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is the name of resource being referenced |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef

**Description**

TypedObjectReference contains enough information to let you locate the typed referenced object

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |
| `namespace` | `string` | Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

**Description**

VolumeResourceRequirements describes the storage resource requirements for a volume.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].fc

**Description**

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `lun` | `integer` | lun is Optional: FC target lun number |
| `readOnly` | `boolean` | readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `targetWWNs` | `array` | targetWWNs is Optional: FC target worldwide names (WWNs) |
| `wwids` | `array` | wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously. |

## .spec.template.spec.volumes[].fc.targetWWNs

**Description**

targetWWNs is Optional: FC target worldwide names (WWNs)

**Type**

`array`

## .spec.template.spec.volumes[].fc.targetWWNs[]

**Type**

`string`

## .spec.template.spec.volumes[].fc.wwids

**Description**

wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

**Type**

`array`

## .spec.template.spec.volumes[].fc.wwids[]

**Type**

`string`

## .spec.template.spec.volumes[].flexVolume

**Description**

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

**Type**

`object`

**Required**

`driver`

| Property | Type | Description |
| --- | --- | --- |
| `driver` | `string` | driver is the name of the driver to use for this volume. |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script. |
| `options` | `object` | options is Optional: this field holds extra command options if any. |
| `readOnly` | `boolean` | readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |

## .spec.template.spec.volumes[].flexVolume.options

**Description**

options is Optional: this field holds extra command options if any.

**Type**

`object`

## .spec.template.spec.volumes[].flexVolume.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: |

| Property | Type | Description |
| --- | --- | --- |
| | | https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].flocker

### Description

Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `datasetName` | `string` | datasetName is Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated |
| `datasetUUID` | `string` | datasetUUID is the UUID of the dataset. This is unique identifier of a Flocker dataset |

## .spec.template.spec.volumes[].gcePersistentDisk

### Description

Represents a Persistent Disk resource in Google Compute Engine. A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

### Type

`object`

### Required

`pdName`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `partition` | `integer` | partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |
| `pdName` | `string` | pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk ↗ |

## .spec.template.spec.volumes[].gitRepo

**Description**

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

**Type**

`object`

**Required**

`repository`

| Property | Type | Description |
| --- | --- | --- |
| `directory` | `string` | directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name. |
| `repository` | `string` | repository is the URL |
| `revision` | `string` | revision is the commit hash for the specified revision. |

## .spec.template.spec.volumes[].glusterfs

**Description**

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`endpoints` `path`

| Property | Type | Description |
| --- | --- | --- |
| `endpoints` | `string` | endpoints is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |
| `path` | `string` | path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `readOnly` | `boolean` | readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod ↗ |

## .spec.template.spec.volumes[].hostPath

### Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

### Type

`object`

### Required

`path`

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗ |
| `type` | `string` | type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath ↗<br><br>Possible enum values:<br>• `""` For backwards compatible, leave it empty if unset<br>• `"BlockDevice"` A block device must exist at the given path<br>• `"CharDevice"` A character device must exist at the given path<br>• `"Directory"` A directory must exist at the given path<br>• `"DirectoryOrCreate"` If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet.<br>• `"File"` A file must exist at the given path<br>• `"FileOrCreate"` If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet.<br>• `"Socket"` A UNIX socket must exist at the given path |

## .spec.template.spec.volumes[].image

### Description

ImageVolumeSource represents a image volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `pullPolicy` | `string` | Policy for pulling OCI objects. Possible values are: Always: the kubelet always attempts to pull the reference. Container creation will fail If the pull fails. Never: the kubelet never pulls the reference and only uses a local image or artifact. |

| Property | Type | Description |
| --- | --- | --- |
| | | Container creation will fail if the reference isn't present. IfNotPresent: the kubelet pulls if the reference isn't already present on disk. Container creation will fail if the reference isn't present and the pull fails. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise.<br><br>Possible enum values:<br><br>- `"Always"` means that kubelet always attempts to pull the latest image. Container will fail If the pull fails.<br>- `"IfNotPresent"` means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails.<br>- `"Never"` means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present |
| `reference` | `string` | Required: Image or artifact reference to be used. Behaves in the same way as pod.spec.containers[*].image. Pull secrets will be assembled in the same way as for the container image by looking up node credentials, SA image pull secrets, and pod spec image pull secrets. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets. |

## .spec.template.spec.volumes[].iscsi

### Description

Represents an ISCSI disk. ISCSI volumes can only be mounted as read/write once. ISCSI volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`targetPortal` `iqn` `lun`

| Property | Type | Description |
| --- | --- | --- |
| `chapAuthDiscovery` | `boolean` | chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication |
| `chapAuthSession` | `boolean` | chapAuthSession defines whether support iSCSI Session CHAP authentication |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi |
| `initiatorName` | `string` | initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection. |
| `iqn` | `string` | iqn is the target iSCSI Qualified Name. |

| Property | Type | Description |
| --- | --- | --- |
| `iscsiInterface` | `string` | iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp). |
| `lun` | `integer` | lun represents iSCSI Target Lun number. |
| `portals` | `array` | portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `targetPortal` | `string` | targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260). |

## .spec.template.spec.volumes[].iscsi.portals

**Description**

portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

**Type**

`array`

## .spec.template.spec.volumes[].iscsi.portals[]

**Type**

`string`

## .spec.template.spec.volumes[].iscsi.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].nfs

**Description**

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`server` `path`

| Property | Type | Description |
|---|---|---|
| `path` | `string` | path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |
| `readOnly` | `boolean` | readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |
| `server` | `string` | server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs |

## .spec.template.spec.volumes[].persistentVolumeClaim

**Description**

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

**Type**

`object`

**Required**

`claimName`

| Property | Type | Description |
|---|---|---|
| `claimName` | `string` | claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims |
| `readOnly` | `boolean` | readOnly Will force the ReadOnly setting in VolumeMounts. Default false. |

## .spec.template.spec.volumes[].photonPersistentDisk

**Description**

Represents a Photon Controller persistent disk resource.

**Type**

`object`

**Required**

`pdID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `pdID` | `string` | pdID is the ID that identifies Photon Controller persistent disk |

## .spec.template.spec.volumes[].portworxVolume

**Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

`object`

**Required**

`volumeID`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `volumeID` | `string` | volumeID uniquely identifies a Portworx volume |

## .spec.template.spec.volumes[].projected

**Description**

Represents a projected volume source

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `sources` | `array` | sources is the list of volume projections. Each entry in this list handles one source. |

## .spec.template.spec.volumes[].projected.sources

**Description**

sources is the list of volume projections. Each entry in this list handles one source.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[]

**Description**

Projection that may be projected along with other supported volume types. Exactly one of these fields must be set.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `clusterTrustBundle` | `object` | ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem. |
| `configMap` | `object` | Adapts a ConfigMap into a projected volume.<br><br>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode. |
| `downwardAPI` | `object` | Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode. |
| `secret` | `object` | Adapts a secret into a projected volume.<br><br>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode. |
| `serviceAccountToken` | `object` | ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise). |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle

**Description**

ClusterTrustBundleProjection describes how to select a set of ClusterTrustBundle objects and project their contents into the pod filesystem.

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `labelSelector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `name` | `string` | Select a single ClusterTrustBundle by object name. Mutually-exclusive with signerName and labelSelector. |
| `optional` | `boolean` | If true, don't block pod startup if the referenced ClusterTrustBundle(s) aren't available. If using name, then the named ClusterTrustBundle is allowed not to exist. If using signerName, then the combination of signerName and labelSelector is allowed to match zero ClusterTrustBundles. |
| `path` | `string` | Relative path from the volume root to write the bundle. |
| `signerName` | `string` | Select all ClusterTrustBundles that match this signer name. Mutually-exclusive with name. The contents of all selected ClusterTrustBundles will be unified and deduplicated. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchExpressions[].values[]

**Type**

`string`

## .spec.template.spec.volumes[].projected.sources[].clusterTrustBundle.labelSelector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.template.spec.volumes[].projected.sources[].configMap

**Description**

Adapts a ConfigMap into a projected volume. The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |
| `optional` | `boolean` | optional specify whether the ConfigMap or its keys must be defined |

## .spec.template.spec.volumes[].projected.sources[].configMap.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].configMap.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI

**Description**

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | Items is a list of DownwardAPIVolume file |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items

**Description**

Items is a list of DownwardAPIVolume file

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[]

**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
| --- | --- | --- |
| `fieldRef` | `object` | ObjectFieldSelector selects an APIVersioned field of an object. |

| Property | Type | Description |
| --- | --- | --- |
| `mode` | `integer` | Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..' |
| `resourceFieldRef` | `object` | ResourceFieldSelector represents container resources (cpu, memory) and their output format |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef

**Description**

ObjectFieldSelector selects an APIVersioned field of an object.

**Type**

`object`

**Required**

`fieldPath`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | Version of the schema the FieldPath is written in terms of, defaults to "v1". |
| `fieldPath` | `string` | Path of the field to select in the specified API version. |

## .spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

`object`

**Required**

`resource`

| Property | Type | Description |
| --- | --- | --- |
| `containerName` | `string` | Container name: required for volumes, optional for env vars |
| `divisor` | `string\|number` | Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and |

| Property | Type | Description |
| --- | --- | --- |
| | | The serialization format is:<br>(Note that <suffix> may be empty, from the "" case in <decimalSI>.)<br><digit> ::= 0 \| 1 \| ... \| 9 <digits> ::= <digit> \| <digit><digits<br>(International System of units; See: http://physics.nist.gov/cuu/Units/binary.ht<br><decimalSI> ::= m \| "" \| k \| M \| G \| T \| P \| E<br>(Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.)<br><decimalExponent> ::= "e" <signedNumber> \| "E" <signedNumber> ```<br>No matter which of the three exponent forms is used, no quantity may represent a num<br>When a Quantity is parsed from a string, it will remember the type of suffix it had,<br>Before serializing, Quantity will be put in "canonical form". This means that Expone<br>- No precision is lost - No fractional digits will be emitted - The exponent (or suf<br>The sign will be omitted unless the number is negative.<br>Examples:<br>- 1.5 will be serialized as "1500m" - 1.5Gi will be serialized as "1536Mi"<br>Note that the quantity will NEVER be internally represented by a floating point numb<br>Non-canonical values will still parse as long as they are well formed, but will be r<br>This format is intended to make it difficult to use these numbers without writing so |
| `resource` | `string` | Required: resource to select |

## .spec.template.spec.volumes[].projected.sources[].secret

### Description

Adapts a secret into a projected volume. The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `items` | `array` | items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names |
| `optional` | `boolean` | optional field specify whether the Secret or its key must be defined |

## .spec.template.spec.volumes[].projected.sources[].secret.items

**Description**

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].projected.sources[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].projected.sources[].serviceAccountToken

**Description**

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

**Type**

`object`

**Required**

`path`

| Property | Type | Description |
|---|---|---|
| `audience` | `string` | audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver. |
| `expirationSeconds` | `integer` | expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours.Defaults to 1 hour and must be at least 10 minutes. |
| `path` | `string` | path is the path relative to the mount point of the file to project the token into. |

## .spec.template.spec.volumes[].quobyte

**Description**

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

**Type**

`object`

**Required**

`registry` `volume`

| Property | Type | Description |
|---|---|---|
| `group` | `string` | group to map volume access to Default is no group |
| `readOnly` | `boolean` | readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false. |
| `registry` | `string` | registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes |
| `tenant` | `string` | tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin |
| `user` | `string` | user to map volume access to Defaults to serivceaccount user |

| Property | Type | Description |
| --- | --- | --- |
| `volume` | `string` | volume is a string that references an already created Quobyte volume by name. |

## .spec.template.spec.volumes[].rbd

### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

### Type

`object`

### Required

`monitors` `image`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd ↗ |
| `image` | `string` | image is the rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `keyring` | `string` | keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `monitors` | `array` | monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `pool` | `string` | pool is the rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `readOnly` | `boolean` | readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `user` | `string` | user is the rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it ↗ |

## .spec.template.spec.volumes[].rbd.monitors

**Description**

monitors is a collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

**Type**

`array`

## .spec.template.spec.volumes[].rbd.monitors[]

**Type**

`string`

## .spec.template.spec.volumes[].rbd.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].scaleIO

**Description**

ScaleIOVolumeSource represents a persistent ScaleIO volume

**Type**

`object`

**Required**

`gateway` `system` `secretRef`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs". |
| `gateway` | `string` | gateway is the host address of the ScaleIO API Gateway. |
| `protectionDomain` | `string` | protectionDomain is the name of the ScaleIO Protection Domain for the configured storage. |
| `readOnly` | `boolean` | readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |

| Property | Type | Description |
| --- | --- | --- |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `sslEnabled` | `boolean` | sslEnabled Flag enable/disable SSL communication with Gateway, default false |
| `storageMode` | `string` | storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned. |
| `storagePool` | `string` | storagePool is the ScaleIO Storage Pool associated with the protection domain. |
| `system` | `string` | system is the name of the storage system as configured in ScaleIO. |
| `volumeName` | `string` | volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source. |

## .spec.template.spec.volumes[].scaleIO.secretRef

**Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].secret

**Description**

Adapts a Secret into a volume. The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `defaultMode` | `integer` | defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON |

| Property | Type | Description |
| --- | --- | --- |
| | | requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |
| `items` | `array` | items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'. |
| `optional` | `boolean` | optional field specify whether the Secret or its keys must be defined |
| `secretName` | `string` | secretName is the name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret ↗ |

## .spec.template.spec.volumes[].secret.items

**Description**

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type**

`array`

## .spec.template.spec.volumes[].secret.items[]

**Description**

Maps a string key to a path within a volume.

**Type**

`object`

**Required**

`key` `path`

| Property | Type | Description |
| --- | --- | --- |
| `key` | `string` | key is the key to project. |
| `mode` | `integer` | mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set. |

| Property | Type | Description |
| --- | --- | --- |
| `path` | `string` | path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'. |

## .spec.template.spec.volumes[].storageos

### Description

Represents a StorageOS persistent volume resource.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `fsType` | `string` | fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `readOnly` | `boolean` | readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. |
| `secretRef` | `object` | LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace. |
| `volumeName` | `string` | volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace. |
| `volumeNamespace` | `string` | volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created. |

## .spec.template.spec.volumes[].storageos.secretRef

### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name of the referent. This field is effectively required, but due to backwards compatibility is allowed to be empty. Instances of this type with an empty value here are almost certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗ |

## .spec.template.spec.volumes[].vsphereVolume

**Description**

Represents a vSphere volume resource.

**Type**

`object`

**Required**

`volumePath`

| Property | Type | Description |
|---|---|---|
| `fsType` | `string` | fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. |
| `storagePolicyID` | `string` | storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName. |
| `storagePolicyName` | `string` | storagePolicyName is the storage Policy Based Management (SPBM) profile name. |
| `volumePath` | `string` | volumePath is the path that identifies vSphere volume vmdk |

## .spec.updateStrategy

**Description**

StatefulSetUpdateStrategy indicates the strategy that the StatefulSet controller will use to perform updates. It includes any additional parameters necessary to perform the update for the indicated strategy.

**Type**

`object`

| Property | Type | Description |
|---|---|---|
| `rollingUpdate` | `object` | RollingUpdateStatefulSetStrategy is used to communicate parameter for RollingUpdateStatefulSetStrategyType. |
| `type` | `string` | Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.<br><br>Possible enum values:<br><br>• `"OnDelete"` triggers the legacy behavior. Version tracking and ordered rolling restarts are disabled. Pods are recreated from the StatefulSetSpec when they are manually deleted. When a scale operation is performed with this strategy,specification version indicated by the StatefulSet's currentRevision.<br>• `"RollingUpdate"` indicates that update will be applied to all Pods in the StatefulSet with respect to the StatefulSet ordering constraints. When a scale operation is performed with this strategy, new Pods will be created from the specification version indicated by the StatefulSet's updateRevision. |

## .spec.updateStrategy.rollingUpdate

**Description**

RollingUpdateStatefulSetStrategy is used to communicate parameter for RollingUpdateStatefulSetStrategyType.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `maxUnavailable` | `integer\|string` | IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number. |
| `partition` | `integer` | Partition indicates the ordinal at which the StatefulSet should be partitioned for updates. During a rolling update, all pods from ordinal Replicas-1 to Partition are updated. All pods from ordinal Partition-1 to 0 remain untouched. This is helpful in being able to do a canary based deployment. The default value is 0. |

## .spec.volumeClaimTemplates

**Description**

volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

**Type**

`array`

## .spec.volumeClaimTemplates[]

**Description**

PersistentVolumeClaim is a user's request for and claim to a persistent volume

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `apiVersion` | `string` | APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources ↗ |
| `kind` | `string` | Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds ↗ |
| `metadata` | `ObjectMeta` | ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create. |

| Property | Type | Description |
| --- | --- | --- |
| `spec` | `object` | PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes |
| `status` | `object` | PersistentVolumeClaimStatus is the current status of a persistent volume claim. |

## .spec.volumeClaimTemplates[].spec

### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `accessModes` | `array` | accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1 ↗ |
| `dataSource` | `object` | TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace. |
| `dataSourceRef` | `object` | TypedObjectReference contains enough information to let you locate the typed referenced object |
| `resources` | `object` | VolumeResourceRequirements describes the storage resource requirements for a volume. |
| `selector` | `object` | A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects. |
| `storageClassName` | `string` | storageClassName is the name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1 ↗ |
| `volumeAttributesClassName` | `string` | volumeAttributesClassName may be used to set the VolumeAttributesClass used by this claim. If specified, the CSI driver will create or update the volume with the attributes defined in the corresponding VolumeAttributesClass. This has a different purpose than storageClassName, it can be changed after the claim is created. An empty string value means that no VolumeAttributesClass will be applied to the claim but it's not allowed to reset this field to empty string once it is set. If unspecified and the PersistentVolumeClaim is unbound, the default VolumeAttributesClass will be set by the persistentvolume controller if it exists. If the resource referred to by volumeAttributesClass does not exist, this PersistentVolumeClaim will be set to a Pending state, as reflected by the modifyVolumeStatus field, until such as a resource exists. More info: |

| Property | Type | Description |
| --- | --- | --- |
| | | https://kubernetes.io/docs/concepts/storage/volume-attributes-classes/ ↗ (Beta) Using this field requires the VolumeAttributesClass feature gate to be enabled (off by default). |
| `volumeMode` | `string` | volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.<br><br>Possible enum values:<br>• `"Block"` means the volume will not be formatted with a filesystem and will remain a raw block device.<br>• `"Filesystem"` means the volume will be or is formatted with a filesystem. |
| `volumeName` | `string` | volumeName is the binding reference to the PersistentVolume backing this claim. |

## .spec.volumeClaimTemplates[].spec.accessModes

**Description**

accessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

**Type**

`array`

## .spec.volumeClaimTemplates[].spec.accessModes[]

**Type**

`string`

## .spec.volumeClaimTemplates[].spec.dataSource

**Description**

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

**Type**

`object`

**Required**

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |

| Property | Type | Description |
| --- | --- | --- |
| `name` | `string` | Name is the name of resource being referenced |

## .spec.volumeClaimTemplates[].spec.dataSourceRef

### Description

TypedObjectReference contains enough information to let you locate the typed referenced object

### Type

`object`

### Required

`kind` `name`

| Property | Type | Description |
| --- | --- | --- |
| `apiGroup` | `string` | APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required. |
| `kind` | `string` | Kind is the type of resource being referenced |
| `name` | `string` | Name is the name of resource being referenced |
| `namespace` | `string` | Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled. |

## .spec.volumeClaimTemplates[].spec.resources

### Description

VolumeResourceRequirements describes the storage resource requirements for a volume.

### Type

`object`

| Property | Type | Description |
| --- | --- | --- |
| `limits` | `object` | Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

| Property | Type | Description |
| --- | --- | --- |
| `requests` | `object` | Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗ |

## .spec.volumeClaimTemplates[].spec.resources.limits

**Description**

Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.volumeClaimTemplates[].spec.resources.requests

**Description**

Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/

**Type**

`object`

## .spec.volumeClaimTemplates[].spec.selector

**Description**

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `matchExpressions` | `array` | matchExpressions is a list of label selector requirements. The requirements are ANDed. |
| `matchLabels` | `object` | matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed. |

## .spec.volumeClaimTemplates[].spec.selector.matchExpressions

**Description**

matchExpressions is a list of label selector requirements. The requirements are ANDed.

**Type**

`array`

## .spec.volumeClaimTemplates[].spec.selector.matchExpressions[]

**Description**

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

**Type**

`object`

**Required**

`key` `operator`

| Property | Type | Description |
|---|---|---|
| `key` | `string` | key is the label key that the selector applies to. |
| `operator` | `string` | operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist. |
| `values` | `array` | values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch. |

## .spec.volumeClaimTemplates[].spec.selector.matchExpressions[].values

**Description**

values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

**Type**

`array`

## .spec.volumeClaimTemplates[].spec.selector.matchExpressions[].values[]

**Type**

`string`

## .spec.volumeClaimTemplates[].spec.selector.matchLabels

**Description**

matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

**Type**

`object`

## .spec.volumeClaimTemplates[].status

**Description**

PersistentVolumeClaimStatus is the current status of a persistent volume claim.

**Type**

`object`

| Property | Type | Description |
| --- | --- | --- |
| `accessModes` | `array` | accessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1 |
| `allocatedResourceStatuses` | `object` | allocatedResourceStatuses stores status of resource being resized for the given PVC. Key names follow standard Kubernetes label syntax. Valid values are either: * Un-prefixed keys: - storage - the capacity of the volume. * Custom resources must use implementation-defined prefixed names such as "example.com/my-custom-resource" Apart from above values - keys that are unprefixed or have kubernetes.io prefix are considered reserved and hence may not be used.<br><br>ClaimResourceStatus can be in any of following states: - ControllerResizeInProgress: State set when resize controller starts resizing the volume in control-plane. - ControllerResizeFailed: State set when resize has failed in resize controller with a terminal error. - NodeResizePending: State set when resize controller has finished resizing the volume but further resizing of volume is needed on the node. - NodeResizeInProgress: State set when kubelet starts resizing the volume. - NodeResizeFailed: State set when resizing has failed in kubelet with a terminal error. Transient errors don't set NodeResizeFailed. For example: if expanding a PVC for more capacity - this field can be one of the following states: - pvc.status.allocatedResourceStatus['storage'] = "ControllerResizeInProgress" - pvc.status.allocatedResourceStatus['storage'] = "ControllerResizeFailed" - pvc.status.allocatedResourceStatus['storage'] = "NodeResizePending" - pvc.status.allocatedResourceStatus['storage'] = "NodeResizeInProgress" - pvc.status.allocatedResourceStatus['storage'] = "NodeResizeFailed" When this field is not set, it means that no resize operation is in progress for the given PVC.<br><br>A controller that receives PVC update with previously unknown resourceName or ClaimResourceStatus should ignore the update for the purpose it was designed. For example - a controller that only is responsible for resizing capacity of the volume, should ignore PVC updates that change other valid resources associated with PVC.<br><br>This is an alpha field and requires enabling RecoverVolumeExpansionFailure feature. |

| Property | Type | Description |
| --- | --- | --- |
| `allocatedResources` | `object` | allocatedResources tracks the resources allocated to a PVC including its capacity. Key names follow standard Kubernetes label syntax. Valid values are either: * Un-prefixed keys: - storage - the capacity of the volume. * Custom resources must use implementation-defined prefixed names such as "example.com/my-custom-resource" Apart from above values - keys that are unprefixed or have kubernetes.io prefix are considered reserved and hence may not be used.<br><br>Capacity reported here may be larger than the actual capacity when a volume expansion operation is requested. For storage quota, the larger value from allocatedResources and PVC.spec.resources is used. If allocatedResources is not set, PVC.spec.resources alone is used for quota calculation. If a volume expansion capacity request is lowered, allocatedResources is only lowered if there are no expansion operations in progress and if the actual volume capacity is equal or lower than the requested capacity.<br><br>A controller that receives PVC update with previously unknown resourceName should ignore the update for the purpose it was designed. For example - a controller that only is responsible for resizing capacity of the volume, should ignore PVC updates that change other valid resources associated with PVC.<br><br>This is an alpha field and requires enabling RecoverVolumeExpansionFailure feature. |
| `capacity` | `object` | capacity represents the actual resources of the underlying volume. |
| `conditions` | `array` | conditions is the current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'Resizing'. |
| `currentVolumeAttributesClassName` | `string` | currentVolumeAttributesClassName is the current name of the VolumeAttributesClass the PVC is using. When unset, there is no VolumeAttributeClass applied to this PersistentVolumeClaim This is a beta field and requires enabling VolumeAttributesClass feature (off by default). |
| `modifyVolumeStatus` | `object` | ModifyVolumeStatus represents the status object of ControllerModifyVolume operation |
| `phase` | `string` | phase represents the current phase of PersistentVolumeClaim.<br><br>Possible enum values:<br>- `"Bound"` used for PersistentVolumeClaims that are bound<br>- `"Lost"` used for PersistentVolumeClaims that lost their underlying PersistentVolume. The claim was bound to a PersistentVolume and this volume does not exist any longer and all data on it was lost.<br>- `"Pending"` used for PersistentVolumeClaims that are not yet bound |

## .spec.volumeClaimTemplates[].status.accessModes

**Description**

accessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

**Type**

`array`

## .spec.volumeClaimTemplates[].status.accessModes[]

**Type**

`string`

## .spec.volumeClaimTemplates[].status.allocatedResourceStatuses

**Description**

allocatedResourceStatuses stores status of resource being resized for the given PVC. Key names follow standard Kubernetes label syntax. Valid values are either: * Un-prefixed keys: - storage - the capacity of the volume. * Custom resources must use implementation-defined prefixed names such as "example.com/my-custom-resource" Apart from above values - keys that are unprefixed or have kubernetes.io prefix are considered reserved and hence may not be used. ClaimResourceStatus can be in any of following states: - ControllerResizeInProgress: State set when resize controller starts resizing the volume in control-plane. - ControllerResizeFailed: State set when resize has failed in resize controller with a terminal error. - NodeResizePending: State set when resize controller has finished resizing the volume but further resizing of volume is needed on the node. - NodeResizeInProgress: State set when kubelet starts resizing the volume. - NodeResizeFailed: State set when resizing has failed in kubelet with a terminal error. Transient errors don't set NodeResizeFailed. For example: if expanding a PVC for more capacity - this field can be one of the following states: - pvc.status.allocatedResourceStatus['storage'] = "ControllerResizeInProgress" - pvc.status.allocatedResourceStatus['storage'] = "ControllerResizeFailed" - pvc.status.allocatedResourceStatus['storage'] = "NodeResizePending" - pvc.status.allocatedResourceStatus['storage'] = "NodeResizeInProgress" - pvc.status.allocatedResourceStatus['storage'] = "NodeResizeFailed" When this field is not set, it means that no resize operation is in progress for the given PVC. A controller that receives PVC update with previously unknown resourceName or ClaimResourceStatus should ignore the update for the purpose it was designed. For example - a controller that only is responsible for resizing capacity of the volume, should ignore PVC updates that change other valid resources associated with PVC. This is an alpha field and requires enabling RecoverVolumeExpansionFailure feature.

**Type**

`object`

## .spec.volumeClaimTemplates[].status.allocatedResources

**Description**

allocatedResources tracks the resources allocated to a PVC including its capacity. Key names follow standard Kubernetes label syntax. Valid values are either: * Un-prefixed keys: - storage - the capacity of the volume. * Custom resources must use implementation-defined prefixed names such as "example.com/my-custom-resource" Apart from above values - keys that are unprefixed or have kubernetes.io prefix are considered reserved and hence may not be used. Capacity reported here may be larger than the actual capacity when a volume expansion operation is requested. For storage quota, the larger value from allocatedResources and PVC.spec.resources is used. If allocatedResources is not set, PVC.spec.resources alone is used for quota calculation. If a volume expansion capacity request is lowered, allocatedResources is only lowered if there are no expansion operations in progress and if the actual volume capacity is equal or lower than the requested capacity. A controller that receives PVC update with previously unknown resourceName should ignore the update for the purpose it was designed. For example - a controller that only is responsible for resizing capacity of the volume, should ignore PVC updates that change other valid resources associated with PVC. This is an alpha field and requires enabling RecoverVolumeExpansionFailure feature.

**Type**

`object`

## .spec.volumeClaimTemplates[].status.capacity

**Description**

capacity represents the actual resources of the underlying volume.

**Type**

`object`

## .spec.volumeClaimTemplates[].status.conditions

**Description**

conditions is the current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'Resizing'.

**Type**

`array`

## .spec.volumeClaimTemplates[].status.conditions[]

**Description**

PersistentVolumeClaimCondition contains details about state of pvc

**Type**

`object`

**Required**

`type` `status`

| Property | Type | Description |
| --- | --- | --- |
| `lastProbeTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `lastTransitionTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | message is the human-readable message indicating details about last transition. |
| `reason` | `string` | reason is a unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "Resizing" that means the underlying persistent volume is being resized. |
| `status` | `string` | Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/reference/kubernetes-api/config-and-storage-resources/persistent-volume-claim-v1/#:~:text=state%20of%20pvc-,conditions.status,-(string)%2C%20required ↗ |
| `type` | `string` | Type is the type of the condition. More info: https://kubernetes.io/docs/reference/kubernetes-api/config-and-storage-resources/persistent-volume-claim-v1/#:~:text=set%20to%20%27ResizeStarted%27.-,PersistentVolumeClaimCondition,-contains%20details%20about ↗ |

## .spec.volumeClaimTemplates[].status.modifyVolumeStatus

**Description**

ModifyVolumeStatus represents the status object of ControllerModifyVolume operation

**Type**

`object`

**Required**

`status`

| Property | Type | Description |
|---|---|---|
| `status` | `string` | status is the status of the ControllerModifyVolume operation. It can be in any of following states:<br>• Pending Pending indicates that the PersistentVolumeClaim cannot be modified due to unmet requirements, such as the specified VolumeAttributesClass not existing.<br>• InProgress InProgress indicates that the volume is being modified.<br>• Infeasible Infeasible indicates that the request has been rejected as invalid by the CSI driver. To resolve the error, a valid VolumeAttributesClass needs to be specified. Note: New statuses can be added in the future. Consumers should check for unknown statuses and fail appropriately.<br>Possible enum values:<br>• `"InProgress"` InProgress indicates that the volume is being modified<br>• `"Infeasible"` Infeasible indicates that the request has been rejected as invalid by the CSI driver. To resolve the error, a valid VolumeAttributesClass needs to be specified<br>• `"Pending"` Pending indicates that the PersistentVolumeClaim cannot be modified due to unmet requirements, such as the specified VolumeAttributesClass not existing |
| `targetVolumeAttributesClassName` | `string` | targetVolumeAttributesClassName is the name of the VolumeAttributesClass the PVC currently being reconciled |

## .status

**Description**

StatefulSetStatus represents the current state of a StatefulSet.

**Type**

`object`

**Required**

`replicas`

| Property | Type | Description |
|---|---|---|
| `availableReplicas` | `integer` | Total number of available pods (ready for at least minReadySeconds) targeted by this statefulset. |
| `collisionCount` | `integer` | collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision. |

| Property | Type | Description |
| --- | --- | --- |
| `conditions` | `array` | Represents the latest available observations of a statefulset's current state. |
| `currentReplicas` | `integer` | currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision. |
| `currentRevision` | `string` | currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas). |
| `observedGeneration` | `integer` | observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server. |
| `readyReplicas` | `integer` | readyReplicas is the number of pods created for this StatefulSet with a Ready Condition. |
| `replicas` | `integer` | replicas is the number of Pods created by the StatefulSet controller. |
| `updateRevision` | `string` | updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas) |
| `updatedReplicas` | `integer` | updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision. |

## .status.conditions

**Description**

Represents the latest available observations of a statefulset's current state.

**Type**

`array`

## .status.conditions[]

**Description**

StatefulSetCondition describes the state of a statefulset at a certain point.

**Type**

`object`

**Required**

`type` `status`

| Property | Type | Description |
| --- | --- | --- |
| `lastTransitionTime` | `string` | Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers. |
| `message` | `string` | A human readable message indicating details about the transition. |
| `reason` | `string` | The reason for the condition's last transition. |
| `status` | `string` | Status of the condition, one of True, False, Unknown. |
| `type` | `string` | Type of statefulset condition. |

# API Endpoints

The following API endpoints are available:

- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/statefulsets`
  - `DELETE`: delete collection of StatefulSet
  - `GET`: list objects of kind StatefulSet
  - `POST`: create a new StatefulSet
- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/statefulsets/{name}`
  - `DELETE`: delete the specified StatefulSet
  - `GET`: read the specified StatefulSet
  - `PATCH`: partially update the specified StatefulSet
  - `PUT`: replace the specified StatefulSet
- `/kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/status`
  - `GET`: read status of the specified StatefulSet
  - `PATCH`: partially update status of the specified StatefulSet
  - `PUT`: replace status of the specified StatefulSet

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/statefulsets

**HTTP method**

`DELETE`

**Description**

delete collection of StatefulSet

**HTTP responses**

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

### Description

list objects of kind StatefulSet

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `StatefulSetList` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`POST`

### Description

create a new StatefulSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `StatefulSet` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `StatefulSet` schema |
| 201 - Created | `StatefulSet` schema |
| 202 - Accepted | `StatefulSet` schema |
| 401 - Unauthorized | Empty |

# /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/statefulsets/{name}

### HTTP method

`DELETE`

### Description

delete the specified StatefulSet

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `Status` schema |
| 202 - Accepted | `Status` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`GET`

### Description

read the specified StatefulSet

### HTTP responses

| HTTP code | Response body |
|---|---|
| 200 - OK | `StatefulSet` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update the specified StatefulSet

### Query parameters

| Parameter | Type | Description |
|---|---|---|
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the |

| Parameter | Type | Description |
| --- | --- | --- |
| | | object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `StatefulSet` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace the specified StatefulSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `StatefulSet` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `StatefulSet` schema |
| 201 - Created | `StatefulSet` schema |
| 401 - Unauthorized | Empty |

## /kubernetes/{cluster}/apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/status

### HTTP method

`GET`

### Description

read status of the specified StatefulSet

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `StatefulSet` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PATCH`

### Description

partially update status of the specified StatefulSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `StatefulSet` schema |
| 401 - Unauthorized | Empty |

### HTTP method

`PUT`

### Description

replace status of the specified StatefulSet

### Query parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `dryRun` | `string` | When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed |
| `fieldValidation` | `string` | fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the |

| Parameter | Type | Description |
| --- | --- | --- |
| | | object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered. |

### Body parameters

| Parameter | Type | Description |
| --- | --- | --- |
| `body` | `StatefulSet` schema | `application/json` formatted |

### HTTP responses

| HTTP code | Response body |
| --- | --- |
| 200 - OK | `StatefulSet` schema |
| 201 - Created | `StatefulSet` schema |
| 401 - Unauthorized | Empty |