



[Alauda Cache Service for Redis OSS](#) > [API 文档](#) > [Kubernetes APIs](#) > Redis APIs

Redis APIs

Redis

Redis 用户



Redis

描述

Redis is the Schema for the redis API

类型

object

规格

属性	类型	描述
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

属性	类型	描述
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	RedisSpec defines the desired state of a Redis instance.
status	object	RedisStatus defines the observed state of a Redis instance.

.spec

描述

RedisSpec defines the desired state of a Redis instance.

类型

object

必填

archresourcesversion

属性	类型	描述
affinity	object	Affinity defines pod scheduling constraints to control pod placement. See: https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#affinity-and-anti-affinity ↗
affinityPolicy	string	AffinityPolicy provides predefined affinity policies for Redis pod placement. Options: SoftAntiAffinity, AntiAffinityInSharding, AntiAffinity Default: SoftAntiAffinity
arch	string	Arch specifies the Redis deployment architecture. Options: cluster, sentinel, standalone
backup	object	Backup configures automated backup settings for Redis data. Controls backup schedule, retention policy, and storage location.
customConfig	object	CustomConfig specifies custom Redis configuration settings that override defaults. These can be set as key-value pairs corresponding to redis.conf parameters.

属性	类型	描述
<code>enableActiveRedis</code>	<code>boolean</code>	EnableActiveRedis activates Redis in active-active replication mode. Required for multi-datacenter or multi-region deployments.
<code>enableTLS</code>	<code>boolean</code>	EnableTLS enables secure TLS connections for Redis. When enabled, Redis will communicate using encrypted connections.
<code>exporter</code>	<code>object</code>	Exporter configures a Redis metrics exporter for Prometheus integration. When enabled, exports Redis metrics in Prometheus format.
<code>expose</code>	<code>object</code>	Expose configures how Redis is exposed outside the Kubernetes cluster. Controls service types, node ports, and load balancer settings.
<code>ipFamilyPrefer</code>	<code>string</code>	IPFamilyPrefer specifies the preferred IP family for Redis networking. Controls whether IPv4 or IPv6 is used for pod and service communication.
<code>nodeSelector</code>	<code>object</code>	NodeSelector constrains Redis pods to nodes with matching labels. See: https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodeselector

属性	类型	描述
<code>password</code>	<code>string</code>	Password is the authentication credential for Redis. Deprecated: Use PasswordSecret instead for improved security.
<code>passwordSecret</code>	<code>string</code>	PasswordSecret references a Kubernetes secret containing the Redis password. The secret must have a 'password' key with a value matching: <code>^[a-zA-Z0-9_!@#\$%^&*()-_+=?]{8,32}\$</code>
<code>patches</code>	<code>object</code>	Patches allows customization of generated Kubernetes resources. Enables fine-grained control over services and other child resources.
<code>pause</code>	<code>boolean</code>	Pause temporarily stops Redis reconciliation when set to true. Useful for maintenance operations or debugging.
<code>persistent</code>	<code>object</code>	Persistent configures persistent storage for Redis data. If specified, a PersistentVolumeClaim will be created for each Redis node.
<code>persistentSize</code>	<code>integer string</code>	PersistentSize specifies the size of the persistent volume for each Redis node. Only applicable when Persistent is configured.

属性	类型	描述
podAnnotations	object	PodAnnotations are custom annotations to add to all Redis pods. Useful for integrations with service meshes, monitoring tools, etc.
redisProxy	object	RedisProxy configures an optional proxy layer in front of Redis. Deprecated: Will be removed in 3.20.
replicas	object	Replicas defines the number of Redis nodes for different architectures. Configuration varies based on the selected architecture (cluster or sentinel).
resources	object	Resources defines the compute resource requirements for Redis containers. This includes CPU and memory limits/requests.
restore	object	Restore configures data restoration from a backup. Specifies the backup source and restoration parameters.
securityContext	object	SecurityContext defines security settings for Redis pods, including user/group IDs, filesystem permissions, and SELinux context.

属性	类型	描述
<code>sentinel</code>	<code>object</code>	Sentinel configures Redis Sentinel for high availability. Only applicable when arch is set to "sentinel".
<code>sentinelCustomConfig</code>	<code>object</code>	SentinelCustomConfig defines custom Redis Sentinel configuration settings. Deprecated: Use sentinel.customConfig instead for improved organization.
<code>serviceID</code>	<code>integer</code>	ServiceID specifies the unique identifier for this Redis instance in an active-active setup. Must be unique across all Active-Active Redis instances.
<code>tolerations</code>	<code>array</code>	Tolerations allow Redis pods to be scheduled on nodes with matching taints. See: https://kubernetes.io/docs/concepts/scheduling-eviction/taint-and-toleration/
<code>upgradeOption</code>	<code>object</code>	UpgradeOption defines the upgrade strategy and automation settings. Controls how and when Redis instances are upgraded.
<code>version</code>	<code>string</code>	Version specifies the Redis version to deploy. Currently supports: 5.0, 6.0, 6.2, 7.0, 7.2, 7.4

.spec.affinity

描述

Affinity defines pod scheduling constraints to control pod placement. See:

<https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#affinity-and-anti-affinity>

类型

object

属性	类型	描述
<code>nodeAffinity</code>	object	Describes node affinity scheduling rules for the pod.
<code>podAffinity</code>	object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
<code>podAntiAffinity</code>	object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

.spec.affinity.nodeAffinity

描述

Describes node affinity scheduling rules for the pod.

类型

object

属性	类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	<code>array</code>	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field. 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, requiredDuringSchedulingIgnoredDuringExecution 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.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	<code>object</code>	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node.

属性	类型	描述
		If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to update), the system may or may not try to eventually evict the pod from its node.

`.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution`

描述

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.

类型

array

`.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]`

描述

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).

类型

object

必填

preference

weight

属性	类型	描述
preference	object	A node selector term, associated with the corresponding weight.
weight	integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference

描述

A node selector term, associated with the corresponding weight.

类型

object

属性	类型	描述
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`

描述

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

类型

array

`.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]`

描述

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

类型

object

必填

keyoperator

属性	类型	描述
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.
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

属性	类型	描述
		element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values

描述

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.

类型

array

.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values[]

类型

string

.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields

描述

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

类型

array

.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.
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

描述

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.

类型

array

.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]

类型

string

.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution

描述

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 an update), the system may or may not try to eventually evict the pod from its node.

类型

object

必填

nodeSelectorTerms

属性	类型	描述
<code>nodeSelectorTerms</code>	array	Required. A list of node selector terms. The terms are ORed.

`.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms`

描述

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

类型

array

`.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]`

描述

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.

类型

object

属性	类型	描述
<code>matchExpressions</code>	array	A list of node selector requirements by node's labels.
<code>matchFields</code>	array	A list of node selector requirements by node's fields.

`.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions`

描述

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

类型

array

.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

描述

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

类型

object

必填

key

operator

属性	类型	描述
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.
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

描述

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.

类型

array

.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]

类型

string

.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields

描述

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

类型

array

.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.
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

描述

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.

类型

array

.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values[]

类型

string

.spec.affinity.podAffinity

描述

Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).

类型

object

属性	类型	描述
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

属性	类型	描述
		scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution, affinity expressions, etc.) to compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has a label 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 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

属性	类型	描述
		intersected, i.e. all terms must be satisfied.

`.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution`

描述

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.

类型

array

`.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]`

描述

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

类型

object

必填

podAffinityTerm

weight

属性	类型	描述
<code>podAffinityTerm</code>	<code>object</code>	Required. A pod affinity term, associated with the corresponding weight.
<code>weight</code>	<code>integer</code>	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

`.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm`

描述

Required. A pod affinity term, associated with the corresponding weight.

类型

`object`

必填

`topologyKey`

属性	类型	描述
<code>labelSelector</code>	<code>object</code>	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
<code>matchLabelKeys</code>	<code>array</code>	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 <code>labelSelector</code> as <code>key in (value)</code> to select

属性	类型	描述
		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).
<code>mismatchLabelKeys</code>	<code>array</code>	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 <code>labelSelector</code> as <code>key notin (value)</code> 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).
<code>namespaceSelector</code>	<code>object</code>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's

属性	类型	描述
		namespace". An empty selector ({}) matches all namespaces.
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

描述

A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.

类型

object

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	<code>object</code>	<code>matchLabels</code> is a map of {key,value} pairs. A single {key,value} in the <code>matchLabels</code> map is equivalent to an element of <code>matchExpressions</code> , 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`

描述

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

类型

`array`

`.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]`

描述

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

类型

`object`

必填

key

operator

属性	类型	描述
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

描述

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.

类型

array

.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.match

Expressions[].values[]

类型

string

.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

描述

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.

类型

object

.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

描述

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).

类型

array

`.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]`

类型

string

`.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys`

描述

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).

类型

array

`.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]`

类型

string

`.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector`

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
<code>matchExpressions</code>	array	matchExpressions is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	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

描述

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

类型

array

.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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

描述

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.

类型

array

.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

类型

string

.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

描述

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.

类型

object

.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

描述

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".

类型

array

`.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]`

类型

string

`.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution`

描述

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.

类型

array

`.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]`

描述

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

类型

object

必填

topologyKey

属性	类型	描述
labelSelector	object	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
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 <code>labelSelector</code> as <code>key in (value)</code> 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 <code>labelSelector</code> as <code>key notin (value)</code> 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.

属性	类型	描述
		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 query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.
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`

描述

A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.

类型

object

属性	类型	描述
<code>matchExpressions</code>	array	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	object	<code>matchLabels</code> is a map of {key,value} pairs. A single {key,value} in the <code>matchLabels</code> map is equivalent to an element of <code>matchExpressions</code> , 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`

描述

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

类型

array

`.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]`

描述

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

类型

object

必填

keyoperator

属性	类型	描述
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.requiredDuringSchedulingIgnore
dDuringExecution[].labelSelector.matchExpressions[].val
ues**

描述

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.

类型

array

**.spec.affinity.podAffinity.requiredDuringSchedulingIgnore
dDuringExecution[].labelSelector.matchExpressions[].val
ues[]**

类型

string

**.spec.affinity.podAffinity.requiredDuringSchedulingIgnore
dDuringExecution[].labelSelector.matchLabels**

描述

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.

类型

object

**.spec.affinity.podAffinity.requiredDuringSchedulingIgnore
dDuringExecution[].matchLabelKeys**

描述

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).

类型

array

.spec.affinity.podAffinity.requiredDuringSchedulingIgnore dDuringExecution[].matchLabelKeys[]

类型

string

.spec.affinity.podAffinity.requiredDuringSchedulingIgnore dDuringExecution[].mismatchLabelKeys

描述

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).

类型

array

.spec.affinity.podAffinity.requiredDuringSchedulingIgnore dDuringExecution[].mismatchLabelKeys[]

类型

string

.spec.affinity.podAffinity.requiredDuringSchedulingIgnore dDuringExecution[].namespaceSelector

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
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.requiredDuringSchedulingIgnore
dDuringExecution[].namespaceSelector.matchExpression
S**

描述

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

类型

array

**.spec.affinity.podAffinity.requiredDuringSchedulingIgnore
dDuringExecution[].namespaceSelector.matchExpression**

s[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.requiredDuringSchedulingIgnore
dDuringExecution[].namespaceSelector.matchExpression
s[].values**

描述

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.

类型

array

**.spec.affinity.podAffinity.requiredDuringSchedulingIgnore
dDuringExecution[].namespaceSelector.matchExpression
s[].values[]**

类型

string

**.spec.affinity.podAffinity.requiredDuringSchedulingIgnore
dDuringExecution[].namespaceSelector.matchLabels**

描述

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.

类型

object

**.spec.affinity.podAffinity.requiredDuringSchedulingIgnore
dDuringExecution[].namespaces**

描述

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".

类型

array

.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

类型

string

.spec.affinity.podAntiAffinity

描述

Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

类型

object

属性	类型	描述
preferredDuringSchedulingIgnoredDuringExecution	array	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expression specified by this field. it may choose a node that violates one or more 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, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expression etc.), compute a sum

属性	类型	描述
		iterating through the elements of this field adding "weight" to the sum if the node has p which matches the corresponding podAffinityTerm; the node(s) with the high sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	<code>array</code>	If the anti-affinity requirements specified by this field are not met at the time of pod scheduling time, the pod will not be scheduled on 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 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 intersected, i.e. all terms must be satisfied.

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution

描述

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.

类型

array

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

描述

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

类型

object

必填

podAffinityTerm weight

属性	类型	描述
podAffinityTerm	object	Required. A pod affinity term, associated with the corresponding weight.

属性	类型	描述
<code>weight</code>	<code>integer</code>	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

`.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm`

描述

Required. A pod affinity term, associated with the corresponding weight.

类型

`object`

必填

`topologyKey`

属性	类型	描述
<code>labelSelector</code>	<code>object</code>	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
<code>matchLabelKeys</code>	<code>array</code>	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 <code>labelSelector</code> as <code>key in (value)</code> 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.

属性	类型	描述
		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).
<code>mismatchLabelKeys</code>	<code>array</code>	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 <code>labelSelector</code> as <code>key notin (value)</code> 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).
<code>namespaceSelector</code>	<code>object</code>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

属性	类型	描述
<code>namespaces</code>	<code>array</code>	<code>namespaces</code> 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 <code>namespaceSelector</code> . null or empty namespaces list and null <code>namespaceSelector</code> means "this pod's namespace".
<code>topologyKey</code>	<code>string</code>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the <code>labelSelector</code> in the specified namespaces, where co-located is defined as running on a node whose value of the label with key <code>topologyKey</code> matches that of any node on which any of the selected pods is running. Empty <code>topologyKey</code> is not allowed.

`.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector`

描述

A label query over a set of resources, in this case pods. If it's null, this `PodAffinityTerm` matches with no Pods.

类型

`object`

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.

属性	类型	描述
<code>matchLabels</code>	<code>object</code>	<code>matchLabels</code> is a map of {key,value} pairs. A single {key,value} in the <code>matchLabels</code> map is equivalent to an element of <code>matchExpressions</code> , 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

描述

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

类型

`array`

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

描述

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

类型

`object`

必填

`key` `operator`

属性	类型	描述
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

描述

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.

类型

array

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

类型

string

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

描述

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.

类型

object

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

描述

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).

类型

array

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

类型

string

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

描述

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).

类型

array

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

类型

string

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
<code>matchExpressions</code>	array	matchExpressions is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	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.namespaceSelector.matchExpressions

描述

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

类型

array

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

描述

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

类型

object

必填

keyoperator

属性	类型	描述
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

描述

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.

类型

array

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

类型

string

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

描述

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.

类型

object

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

描述

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".

类型

array

.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

类型

string

.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution

描述

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.

类型

array

.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

描述

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

类型

object

必填

topologyKey

属性	类型	描述
labelSelector	object	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
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 <code>labelSelector</code> as <code>key in (value)</code> 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 <code>labelSelector</code> as <code>key notin (value)</code> 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.

属性	类型	描述
		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 query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.
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`

描述

A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.

类型

object

属性	类型	描述
<code>matchExpressions</code>	array	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	object	<code>matchLabels</code> is a map of {key,value} pairs. A single {key,value} in the <code>matchLabels</code> map is equivalent to an element of <code>matchExpressions</code> , 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`

描述

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

类型

array

`.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]`

描述

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

类型

object

必填

key operator

属性	类型	描述
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

描述

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.

类型

array

.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

类型

string

.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

描述

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.

类型

object

.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

描述

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).

类型

array

`.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]`

类型

string

`.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys`

描述

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).

类型

array

`.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]`

类型

string

`.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector`

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
<code>matchExpressions</code>	array	matchExpressions is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	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`

描述

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

类型

array

`.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpres`

sions[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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[].namespaceSelector.matchExpressions[].values

描述

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.

类型

array

.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

类型

string

.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

描述

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.

类型

object

.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

描述

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".

类型

array

.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

类型

string

.spec.backup

描述

Backup configures automated backup settings for Redis data. Controls backup schedule, retention policy, and storage location.

类型

object

属性	类型	描述
image	string	
schedule	array	

.spec.backup.schedule

类型

array

.spec.backup.schedule[]

类型

object

必填

keep schedule storage

属性	类型	描述
keep	integer	
keepAfterDeletion	boolean	
name	string	
schedule	string	
storage	object	
target	object	

.spec.backup.schedule[].storage

类型

object

属性	类型	描述
size	integer string	
storageClassName	string	

.spec.backup.schedule[].target

类型

object

属性	类型	描述
s3option	object	S3Option

.spec.backup.schedule[].target.s3Option

描述

S3Option

类型

object

属性	类型	描述
bucket	string	
dir	string	
s3Secret	string	

.spec.customConfig

描述

CustomConfig specifies custom Redis configuration settings that override defaults. These can be set as key-value pairs corresponding to redis.conf parameters.

类型

object

.spec.exporter

描述

Exporter configures a Redis metrics exporter for Prometheus integration. When enabled, exports Redis metrics in Prometheus format.

类型

object

属性	类型	描述
enabled	boolean	

属性	类型	描述
<code>image</code>	<code>string</code>	
<code>imagePullPolicy</code>	<code>string</code>	PullPolicy describes a policy for if/when to pull a container image
<code>resources</code>	<code>object</code>	ResourceRequirements describes the compute resource requirements.

.spec.exporter.resources

描述

ResourceRequirements describes the compute resource requirements.

类型

`object`

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

属性	类型	描述
<code>limits</code>	<code>object</code>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗
<code>requests</code>	<code>object</code>	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.exporter.resources.claims`

描述

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.

类型

`array`

`.spec.exporter.resources.claims[]`

描述

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

类型

`object`

必填

`name`

属性	类型	描述
<code>name</code>	<code>string</code>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.
<code>request</code>	<code>string</code>	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.exporter.resources.limits`

描述

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

类型

`object`

`.spec.exporter.resources.requests`

描述

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/>

类型

`object`

`.spec.expose`

描述

Expose configures how Redis is exposed outside the Kubernetes cluster. Controls service types, node ports, and load balancer settings.

类型

object

属性	类型	描述
accessPort	integer	AccessPort defines the lb access nodeport
annotations	object	The annnotations of the service which attached to services
dataStorageNodePortMap	object	NodePortMap defines the map of the nodeport for redis sentinel only Reversed for 3.14 backward compatibility
dataStorageNodePortSequence	string	NodePortMap defines the map of the nodeport for redis nodes NodePortSequence defines the sequence of the nodeport for redis cluster only
enableNodePort	boolean	EnableNodePort defines if the nodeport is enabled

属性	类型	描述
<code>type</code>	<code>string</code>	ServiceType defines the type of the all related service

.spec.expose.annotations

描述

The annotations of the service which attached to services

类型

`object`

.spec.expose.dataStorageNodePortMap

描述

NodePortMap defines the map of the nodeport for redis sentinel only Reversed for 3.14 backward compatibility

类型

`object`

.spec.nodeSelector

描述

NodeSelector constrains Redis pods to nodes with matching labels. See: <https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodeselector>

类型

`object`

.spec.patches

描述

Patches allows customization of generated Kubernetes resources. Enables fine-grained control over services and other child resources.

类型

object

属性	类型	描述
<code>services</code>	<code>array</code>	Services allows customizing the Kubernetes Services created for Redis. Can modify properties like port configuration, load balancer settings, and annotations.

`.spec.patches.services`

描述

Services allows customizing the Kubernetes Services created for Redis. Can modify properties like port configuration, load balancer settings, and annotations.

类型

array

`.spec.patches.services[]`

描述

Patch configuration for the Service created to serve traffic to the cluster. Allows for the manifest of the created Service to be overwritten with custom configuration.

类型

object

属性	类型	描述
metadata	ObjectMeta ↗	Metadata contains Service metadata like name, namespace, labels, and annotations.
spec	object	Spec defines the Service behavior including type, ports, and selectors. See: https://kubernetes.io/docs/concepts/services-networking/service/ ↗

.spec.patches.services[].spec

描述

Spec defines the Service behavior including type, ports, and selectors. See: <https://kubernetes.io/docs/concepts/services-networking/service/>

类型

object

属性	类型	描述
<code>allocateLoadBalancerNodePorts</code>	<code>boolean</code>	<code>allocateLoadBalancerNodePorts</code> defines if NodePorts will be automatically allocated for services with type LoadBalancer. Default is "true". It may be set to "false" if the cluster load-balancer does not rely on NodePorts. If the caller requests specific NodePorts (by specifying a value), those requests will be respected, regardless of this field. This field may only be set for services with type LoadBalancer and will be cleared if the type is changed to any other type.
<code>clusterIP</code>	<code>string</code>	<code>clusterIP</code> is the IP address of the service and is usually assigned randomly. If an address is specified manually, it must be in-range (as per system configuration), and is not in use, it will be allocated to the service; otherwise creation of the service will fail. This field must not be changed through updates unless the type field is also being changed to ExternalName (which requires this field to be blank) or the type field is being changed from ExternalName (in which case this field may optionally be specified, as described above). Valid values are "None", empty string (""), or a valid IP address. Setting this to "None" makes a "headless service" (no virtual IP), which is useful when direct endpoint connections are preferred and proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. If this field is specified when creating a Service of type ExternalName,

属性	类型	描述
		creation will fail. This field will be wiped wh updating a Service to type ExternalName. More info: https://kubernetes.io/docs/concepts/service-networking/service/#virtual-ips-and-service-proxies ↗
clusterIPs	array	ClusterIPs is a list of IP addresses assigne to this service, and are usually assigned randomly. If an address is specified manually, is in-range (as per system configuration), and is not in use, it will be allocated to the service; otherwise creation of the service will fail. This field may not be changed through updates unless the type field is also being changed to ExternalNam (which requires this field to be empty) or th type field is being changed from ExternalName (in which case this field may optionally be specified, as describe above) Valid values are "None", empty string (""), , a valid IP address. Setting this to "None" makes a "headless service" (no virtual IP), which is useful when direct endpoint connections are preferred and proxying is not required. Only applies to types Clusterl NodePort, and LoadBalancer. If this field is specified when creating a Service of type ExternalName, creation will fail. This field v be wiped when updating a Service to type ExternalName. If this field is not specified, will be initialized from the clusterIP field. If this field is specified, clients must ensure

属性	类型	描述
		that clusterIPs[0] and clusterIP have the same value. This field may hold a maximum of two entries (dual-stack IPs, in either order). These IPs must correspond to the values c the ipFamilies field. Both clusterIPs and ipFamilies are governed by the ipFamilyPolicy field. More info: https://kubernetes.io/docs/concepts/service-networking/service/#virtual-ips-and-service-proxies
externalIPs	array	externalIPs is a list of IP addresses for whi nodes in the cluster will also accept traffic i this service. These IPs are not managed b Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load balancers that are not part of the Kubernet system.
externalName	string	externalName is the external reference tha discovery mechanisms will return as an ali for this service (e.g. a DNS CNAME record No proxying will be involved. Must be a lowercase RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires type to be "ExternalName".

属性	类型	描述
<code>externalTrafficPolicy</code>	<code>string</code>	<code>externalTrafficPolicy</code> describes how nodes distribute service traffic they receive on one of the Service's "externally-facing" addresses (NodePorts, ExternalIPs, and LoadBalancer IPs). If set to "Local", the proxy will configure the service in a way that assumes that external load balancers will take care of balancing the service traffic between nodes, and so each node will deliver traffic only to the node-local endpoints of the service, without masquerading the client source IP. (Traffic mistakenly sent to a node with no endpoint will be dropped.) The default value, "Cluster", uses the standard behavior of routing to all endpoints evenly (possibly modified by topology and other features). Note that traffic sent to an External IP or LoadBalancer IP from within the cluster will always get "Cluster" semantics, but clients sending to a NodePort from within the cluster may need to take traffic policy into account when picking a node.
<code>healthCheckNodePort</code>	<code>integer</code>	<code>healthCheckNodePort</code> specifies the healthcheck nodePort for the service. This only applies when type is set to LoadBalancer and <code>externalTrafficPolicy</code> is set to Local. If a value is specified, is in-range, and is not in use, it will be used. If not specified, a value will be automatically allocated. External systems (e.g. load-balancers) can use this port to determine if

属性	类型	描述
		given node holds endpoints for this service or not. If this field is specified when creating a Service which does not need it, creation will fail. This field will be wiped when updating a Service to no longer need it (e.g. changing type). This field cannot be updated once set.
<code>internalTrafficPolicy</code>	<code>string</code>	InternalTrafficPolicy describes how nodes distribute service traffic they receive on the ClusterIP. If set to "Local", the proxy will assume that pods only want to talk to endpoints of the service on the same node as the pod, dropping the traffic if there are local endpoints. The default value, "Cluster", uses the standard behavior of routing to all endpoints evenly (possibly modified by topology and other features).
<code>ipFamilies</code>	<code>array</code>	IPFamilies is a list of IP families (e.g. IPv4, IPv6) assigned to this service. This field is usually assigned automatically based on cluster configuration and the ipFamilyPolicy field. If this field is specified manually, the requested family is available in the cluster, and ipFamilyPolicy allows it, it will be used otherwise creation of the service will fail. This field is conditionally mutable: it allows for adding or removing a secondary IP family, but it does not allow changing the primary IP family of the Service. Valid values are "IPv4" and "IPv6". This field only applies

属性	类型	描述
		to Services of types ClusterIP, NodePort, and LoadBalancer, and does apply to "headless" services. This field will be wiped when updating a Service to type ExternalName. This field may hold a maximum of two entries (dual-stack families, in either order). These families must correspond to the values of the clusterIPs field, if specified. Both clusterIPs and ipFamilies are governed by the ipFamilyPolicy field.
ipFamilyPolicy	string	IPFamilyPolicy represents the dual-stack-ness requested or required by this Service: if there is no value provided, then this field will be set to SingleStack. Services can be "SingleStack" (a single IP family), "PreferDualStack" (two IP families on dual-stack configured clusters or a single IP family on single-stack clusters), or "RequireDualStack" (two IP families on dual-stack configured clusters, otherwise fail). The ipFamilies and clusterIPs fields depend on the value of this field. This field will be wiped when updating a service to type ExternalName.
loadBalancerClass	string	loadBalancerClass is the class of the load balancer implementation this Service belongs to. If specified, the value of this field must be a label-style identifier, with an optional prefix, e.g. "internal-vip" or

属性	类型	描述
		"example.com/internal-vip". Unprefixed names are reserved for end-users. This field can only be set when the Service type is 'LoadBalancer'. If not set, the default load balancer implementation is used, today this is typically done through the cloud provider integration, but should apply for any default implementation. If set, it is assumed that a load balancer implementation is watching for Services with a matching class. Any default load balancer implementation (e.g. cloud providers) should ignore Services that set this field. This field can only be set when creating or updating a Service to type 'LoadBalancer'. Once set, it can not be changed. This field will be wiped when a service is updated to a non 'LoadBalancer' type.
loadBalancerIP	string	Only applies to Service Type: LoadBalancer This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature. Deprecated: This field was underspecified and its meaning varies across implementations. Using it is non-portable and it may not support dual-stack. Users are encouraged to use implementation-specific annotations when available.

属性	类型	描述
loadBalancerSourceRanges	array	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/create-external-load-balancer/ ↗
ports	array	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/service-networking/service/#virtual-ips-and-service-proxies ↗
publishNotReadyAddresses	boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or

属性	类型	描述
		EndpointSlice resources can safely assum this behavior.
selector	object	Route service traffic to pods with label key: and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modifi Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/service networking/service/
sessionAffinity	string	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/service networking/service/#virtual-ips-and-service proxies
sessionAffinityConfig	object	sessionAffinityConfig contains the configurations of session affinity.
trafficDistribution	string	TrafficDistribution offers a way to express preferences for how traffic is distributed to Service endpoints. Implementations can us this field as a hint, but are not required to guarantee strict adherence. If the field is n

属性	类型	描述
		set, the implementation will apply its default routing strategy. If set to "PreferClose", implementations should prioritize endpoints that are topologically close (e.g., same zone). This is a beta field and requires enabling ServiceTrafficDistribution feature.
type	string	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object or EndpointSlice objects. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a virtual IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the same endpoints as the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the same endpoints as the clusterIP. "ExternalName" aliases this service to the specified externalName. Several other fields do not apply to ExternalName services. More info: https://kubernetes.io/docs/concepts/service-networking/service/#publishing-services-service-types

`.spec.patches.services[].spec.clusterIPs`

描述

ClusterIPs is a list of IP addresses assigned to this service, and are usually assigned randomly. If an address is specified manually, is in-range (as per system configuration), and is not in use, it will be allocated to the service; otherwise creation of the service will fail. This field may not be changed through updates unless the type field is also being changed to ExternalName (which requires this field to be empty) or the type field is being changed from ExternalName (in which case this field may optionally be specified, as describe above). Valid values are "None", empty string (""), or a valid IP address. Setting this to "None" makes a "headless service" (no virtual IP), which is useful when direct endpoint connections are preferred and proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. If this field is specified when creating a Service of type ExternalName, creation will fail. This field will be wiped when updating a Service to type ExternalName. If this field is not specified, it will be initialized from the clusterIP field. If this field is specified, clients must ensure that clusterIPs[0] and clusterIP have the same value. This field may hold a maximum of two entries (dual-stack IPs, in either order). These IPs must correspond to the values of the ipFamilies field. Both clusterIPs and ipFamilies are governed by the ipFamilyPolicy field. More info: <https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies>

类型

array

`.spec.patches.services[].spec.clusterIPs[]`

类型

string

`.spec.patches.services[].spec.externalIPs`

描述

externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.

类型

array

.spec.patches.services[].spec.externalIPs[]

类型

string

.spec.patches.services[].spec.ipFamilies

描述

IPFamilies is a list of IP families (e.g. IPv4, IPv6) assigned to this service. This field is usually assigned automatically based on cluster configuration and the ipFamilyPolicy field. If this field is specified manually, the requested family is available in the cluster, and ipFamilyPolicy allows it, it will be used; otherwise creation of the service will fail. This field is conditionally mutable: it allows for adding or removing a secondary IP family, but it does not allow changing the primary IP family of the Service. Valid values are "IPv4" and "IPv6". This field only applies to Services of types ClusterIP, NodePort, and LoadBalancer, and does not apply to "headless" services. This field will be wiped when updating a Service to type ExternalName. This field may hold a maximum of two entries (dual-stack families, in either order). These families must correspond to the values of the clusterIPs field, if specified. Both clusterIPs and ipFamilies are governed by the ipFamilyPolicy field.

类型

array

.spec.patches.services[].spec.ipFamilies[]

描述

IPFamily represents the IP Family (IPv4 or IPv6). This type is used to express the family of an IP expressed by a type (e.g. service.spec.ipFamilies).

类型

string

`.spec.patches.services[].spec.loadBalancerSourceRanges`

描述

If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info:

<https://kubernetes.io/docs/tasks/access-application-cluster/create-external-load-balancer/>

类型

array

`.spec.patches.services[].spec.loadBalancerSourceRanges`

`[]`

类型

string

`.spec.patches.services[].spec.ports`

描述

The list of ports that are exposed by this service. More info:

<https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies>

类型

array

`.spec.patches.services[].spec.ports[]`

描述

ServicePort contains information on service's port.

类型

object

必填

port

属性	类型	描述
appProtocol	string	The application protocol for this port. This is used as a hint for implementations to offer richer behavior for protocols that they understand. This field follows standard Kubernetes label syntax. Valid values are either: • Un-prefixed protocol names - reserved for IANA standard service names (as per RFC-6335 and https://www.iana.org/assignments/service-names). • Kubernetes-defined prefixed names: • 'kubernetes.io/h2c' - HTTP/2 prior knowledge over cleartext as described in https://www.rfc-editor.org/rfc/rfc9113.html#name-starting-http-2-with-prior- • 'kubernetes.io/ws' - WebSocket over cleartext as described in https://www.rfc-editor.org/rfc/rfc6455 • 'kubernetes.io/wss' - WebSocket over TLS as described in https://www.rfc-editor.org/rfc/rfc6455 • Other protocols should use implementation-defined prefixed names such as mycompany.com/my-custom-protocol.
name	string	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the

属性	类型	描述
		EndpointPort. Optional if only one ServicePort is defined on this service.
nodePort	integer	The port on each node on which this service is exposed when type is NodePort or LoadBalancer. Usually assigned by the system. If a value is specified, in-range, and not in use it will be used, otherwise the operation will fail. If not specified, a port will be allocated if this Service requires one. If this field is specified when creating a Service which does not need it, creation will fail. This field will be wiped when updating a Service to no longer need it (e.g. changing type from NodePort to ClusterIP). More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	integer	The port that will be exposed by this service.
protocol	string	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.
targetPort	integer string	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More

属性	类型	描述
		info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service ↗

`.spec.patches.services[].spec.selector`

描述

Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info:

<https://kubernetes.io/docs/concepts/services-networking/service/>

类型

object

`.spec.patches.services[].spec.sessionAffinityConfig`

描述

sessionAffinityConfig contains the configurations of session affinity.

类型

object

属性	类型	描述
<code>clientIP</code>	object	clientIP contains the configurations of Client IP based session affinity.

`.spec.patches.services[].spec.sessionAffinityConfig.clientIP`

描述

clientIP contains the configurations of Client IP based session affinity.

类型

object

属性	类型	描述
timeoutSeconds	integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

.spec.persistent

描述

Persistent configures persistent storage for Redis data. If specified, a PersistentVolumeClaim will be created for each Redis node.

类型

object

必填

storageClassName

属性	类型	描述
storageClassName	string	StorageClassName specifies the name of the Kubernetes StorageClass to use for Redis persistent volumes. The storage class determines provisioner, reclaim policy, and mount options.

.spec.podAnnotations

描述

PodAnnotations are custom annotations to add to all Redis pods. Useful for integrations with service meshes, monitoring tools, etc.

类型

object

.spec.redisProxy

描述

RedisProxy configures an optional proxy layer in front of Redis. Deprecated: Will be removed in 3.20.

类型

object

属性	类型	描述
affinity	object	Affinity defines pod scheduling constraints for proxy pods. Controls pod placement based on node affinity, pod affinity, and anti-affinity.
config	object	Config provides additional configuration parameters for the Redis Proxy. Key-value pairs corresponding to proxy configuration options.
enable	boolean	Enable toggles whether the Redis Proxy service is deployed. When true, a proxy layer will be created between clients and Redis.

属性	类型	描述
<code>image</code>	<code>string</code>	Image specifies the container image used for the Redis Proxy. Format: [registry]/[repository]:[tag]
<code>nodeSelector</code>	<code>object</code>	NodeSelector constrains proxy pods to nodes with matching labels. See: https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodeselector
<code>replicas</code>	<code>integer</code>	Replicas defines how many proxy instances to deploy for load balancing. Higher values provide better availability and throughput.
<code>resources</code>	<code>object</code>	Resources defines compute resource requirements for the proxy containers. Controls CPU and memory limits/requests.
<code>tolerations</code>	<code>array</code>	Tolerations allow the proxy pods to be scheduled on nodes with matching taints. See: https://kubernetes.io/docs/concepts/scheduling-eviction/taint-and-toleration/

.spec.redisProxy.affinity

描述

Affinity defines pod scheduling constraints for proxy pods. Controls pod placement based on node affinity, pod affinity, and anti-affinity.

类型

object

属性	类型	描述
nodeAffinity	object	Describes node affinity scheduling rules for the pod.
podAffinity	object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

.spec.redisProxy.affinity.nodeAffinity

描述

Describes node affinity scheduling rules for the pod.

类型

object

属性	类型	描述
preferredDuringSchedulingIgnoredDuringExecution	array	The scheduler will pri to schedule pods to nodes that satisfy the affinity expressions specified by this field. it may choose a node violates one or more the expressions. The

属性	类型	描述
		node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

属性	类型	描述
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	<code>object</code>	If the affinity requirement specified by this field is not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirement specified by this field ceases to be met at some point during pod execution (e.g. due to a change in resources or a system restart), the system may or may not try to eventually evict the pod from its node.

`.spec.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution`

描述

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.

类型

`array`

`.spec.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]`

描述

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).

类型

`object`

必填

`preference``weight`

属性	类型	描述
<code>preference</code>	<code>object</code>	A node selector term, associated with the corresponding weight.
<code>weight</code>	<code>integer</code>	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

`.spec.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference`

描述

A node selector term, associated with the corresponding weight.

类型

`object`

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	A list of node selector requirements by node's labels.
<code>matchFields</code>	<code>array</code>	A list of node selector requirements by node's fields.

`.spec.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions`

描述

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

类型

`array`

`.spec.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]`

描述

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

类型

`object`

必填

`key`

`operator`

属性	类型	描述
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.
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.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values

描述

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.

类型

array

.spec.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpres

sions[].values[]

类型

string

.spec.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields

描述

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

类型

array

.spec.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

描述

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

类型

object

必填

key

operator

属性	类型	描述
key	string	The label key that the selector applies to.

属性	类型	描述
<code>operator</code>	<code>string</code>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
<code>values</code>	<code>array</code>	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.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values

描述

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.

类型

`array`

.spec.redisProxy.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]

类型

`string`

`.spec.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution`

描述

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 an update), the system may or may not try to eventually evict the pod from its node.

类型

`object`

必填

`nodeSelectorTerms`

属性	类型	描述
<code>nodeSelectorTerms</code>	<code>array</code>	Required. A list of node selector terms. The terms are ORed.

`.spec.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms`

描述

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

类型

`array`

`.spec.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]`

描述

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.

类型

object

属性	类型	描述
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.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions

描述

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

类型

array

.spec.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

描述

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

类型

object

必填

key

operator

属性	类型	描述
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.
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.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values

描述

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.

类型

array

.spec.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]

类型

string

.spec.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields

描述

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

类型

array

.spec.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

描述

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

类型

object

必填

key

operator

属性	类型	描述
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.
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.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

描述

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.

类型

array

.spec.redisProxy.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

hFields[].values[]

类型

string

.spec.redisProxy.affinity.podAffinity

描述

Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).

类型

object

属性	类型	描述
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 v the greatest sum of weights, i.e. for each node that meets all o scheduling requireme (resource request, requiredDuringSched affinity expressions, € compute a sum by iterating through the elements of this field

属性	类型	描述
		adding "weight" to the sum if the node has p which matches the corresponding podAffinityTerm; the node(s) with the high sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	<code>array</code>	If the affinity requirement specified by this field is not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirement specified by this field ceases to be met at some point during pod execution (e.g. due to 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 intersected, i.e. all terms must be satisfied.

`.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution`

描述

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.

类型

`array`

`.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]`

描述

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

类型

`object`

必填

`podAffinityTerm``weight`

属性	类型	描述
<code>podAffinityTerm</code>	<code>object</code>	Required. A pod affinity term, associated with the corresponding weight.

属性	类型	描述
<code>weight</code>	<code>integer</code>	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

`.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm`

描述

Required. A pod affinity term, associated with the corresponding weight.

类型

`object`

必填

`topologyKey`

属性	类型	描述
<code>labelSelector</code>	<code>object</code>	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
<code>matchLabelKeys</code>	<code>array</code>	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 <code>labelSelector</code> as <code>key in (value)</code> 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.

属性	类型	描述
		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).
<code>mismatchLabelKeys</code>	<code>array</code>	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 <code>labelSelector</code> as <code>key notin (value)</code> 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).
<code>namespaceSelector</code>	<code>object</code>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

属性	类型	描述
<code>namespaces</code>	<code>array</code>	<code>namespaces</code> 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 <code>namespaceSelector</code> . null or empty namespaces list and null <code>namespaceSelector</code> means "this pod's namespace".
<code>topologyKey</code>	<code>string</code>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the <code>labelSelector</code> in the specified namespaces, where co-located is defined as running on a node whose value of the label with key <code>topologyKey</code> matches that of any node on which any of the selected pods is running. Empty <code>topologyKey</code> is not allowed.

`.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector`

描述

A label query over a set of resources, in this case pods. If it's null, this `PodAffinityTerm` matches with no Pods.

类型

`object`

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	<code>object</code>	<code>matchLabels</code> is a map of {key,value} pairs. A single {key,value} in the <code>matchLabels</code> map is equivalent to an element of <code>matchExpressions</code> , whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

`.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions`

描述

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

类型

`array`

`.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]`

描述

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

类型

`object`

必填

属性	类型	描述
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.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

描述

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.

类型

array

.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

ctor.matchExpressions[].values[]

类型

string

.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

描述

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.

类型

object

.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

描述

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).

类型

array

`.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]`

类型

string

`.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys`

描述

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).

类型

array

`.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]`

类型

string

.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
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.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions

描述

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

类型

array

.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespace

ceSelector.matchExpressions[].values

描述

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.

类型

array

.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

类型

string

.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

描述

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.

类型

object

.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

描述

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".

类型

array

.spec.redisProxy.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

类型

string

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution

描述

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.

类型

array

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

描述

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

类型

object

必填

topologyKey

属性	类型	描述
<code>labelSelector</code>	object	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
<code>matchLabelKeys</code>	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 <code>labelSelector</code> as <code>key in (value)</code> 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 <code>matchLabelKeys</code> and <code>labelSelector</code> . Also, <code>matchLabelKeys</code> cannot be set when <code>labelSelector</code> isn't set. This is a beta field and requires enabling <code>MatchLabelKeysInPodAffinity</code> feature gate (enabled by default).

属性	类型	描述
<code>mismatchLabelKeys</code>	<code>array</code>	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 <code>labelSelector</code> as <code>key notin (value)</code> 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).
<code>namespaceSelector</code>	<code>object</code>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.
<code>namespaces</code>	<code>array</code>	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".

属性	类型	描述
<code>topologyKey</code>	<code>string</code>	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.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector`

描述

A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.

类型

`object`

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	matchExpressions is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	<code>object</code>	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.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

描述

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

类型

array

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.

属性	类型	描述
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.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

描述

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.

类型

array

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

类型

string

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

描述

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.

类型

object

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

描述

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).

类型

array

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

类型

string

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

描述

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).

类型

array

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

类型

string

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	matchExpressions is a list of label selector requirements. The requirements are ANDed.

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

`.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions`

描述

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

类型

`array`

`.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]`

描述

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

类型

`object`

必填

`key` `operator`

属性	类型	描述
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.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

描述

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.

类型

array

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

类型

`string`

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

描述

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.

类型

`object`

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

描述

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".

类型

`array`

.spec.redisProxy.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

类型

`string`

.spec.redisProxy.affinity.podAntiAffinity

描述

Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

类型

object

属性	类型	描述
preferredDuringSchedulingIgnoredDuringExecution	array	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expression specified by this field, it may choose a node that violates one or more 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, requiredDuringSchedulingIgnoredDuringExecution, anti-affinity expression etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has a label which matches the corresponding podAffinityTerm; the node(s) with the highest sum is the most preferred.

属性	类型	描述
		sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	<code>array</code>	If the anti-affinity requirements specified by this field are not met at the time of pod scheduling time, the pod will not be scheduled to 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 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 intersected, i.e. all terms must be satisfied.

`.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution`

描述

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.

类型

array

.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

描述

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

类型

object

必填

podAffinityTerm

weight

属性	类型	描述
podAffinityTerm	object	Required. A pod affinity term, associated with the corresponding weight.
weight	integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

描述

Required. A pod affinity term, associated with the corresponding weight.

类型

object

必填

topologyKey

属性	类型	描述
labelSelector	object	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
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

属性	类型	描述
		incoming pod labels, those key-value labels are merged with <code>labelSelector</code> as <code>key notin (value)</code> 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 <code>mismatchLabelKeys</code> and <code>labelSelector</code>. Also, <code>mismatchLabelKeys</code> cannot be set when <code>labelSelector</code> isn't set. This is a beta field and requires enabling <code>MatchLabelKeysInPodAffinity</code> feature gate (enabled by default).
<code>namespaceSelector</code>	<code>object</code>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the <code>namespaces</code> field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector (<code>{}</code>) matches all namespaces.
<code>namespaces</code>	<code>array</code>	<code>namespaces</code> 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 <code>namespaceSelector</code>. null or empty namespaces list and null <code>namespaceSelector</code> means "this pod's namespace".
<code>topologyKey</code>	<code>string</code>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the <code>labelSelector</code> in the specified namespaces, where

属性	类型	描述
		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.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector`

描述

A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.

类型

object

属性	类型	描述
<code>matchExpressions</code>	array	matchExpressions is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	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.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.label`

Selector.matchExpressions

描述

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

类型

array

.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.

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

`.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values`

描述

`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.

类型

`array`

`.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]`

类型

`string`

`.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels`

描述

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.

类型

object

.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

描述

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).

类型

array

.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

类型

string

.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mis

matchLabelKeys

描述

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).

类型

array

`.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]`

类型

string

`.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector`

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	<code>object</code>	<code>matchLabels</code> is a map of {key,value} pairs. A single {key,value} in the <code>matchLabels</code> map is equivalent to an element of <code>matchExpressions</code> , whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

`.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces.selector.matchExpressions`

描述

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

类型

`array`

`.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces.selector.matchExpressions[]`

描述

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

类型

`object`

必填

key

operator

属性	类型	描述
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.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces.selector.matchExpressions[].values

描述

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.

类型

array

.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces.selector.matchExpressions[].values

namespaceSelector.matchExpressions[].values[]

类型

string

.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

描述

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.

类型

object

.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

描述

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".

类型

array

.spec.redisProxy.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

类型

string

`.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution`

描述

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.

类型

array

`.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]`

描述

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

类型

object

必填

topologyKey

属性	类型	描述
labelSelector	object	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
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 not 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 mismatchLabelKeys and labelSelector.

属性	类型	描述
		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 query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.
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.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector`

描述

A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.

类型

object

属性	类型	描述
<code>matchExpressions</code>	array	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	object	<code>matchLabels</code> 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.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions`

描述

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

类型

array

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

描述

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.

类型

array

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

类型

string

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

描述

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.

类型

object

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

描述

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).

类型

array

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

类型

string

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys

描述

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).

类型

array

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]

类型

string

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
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.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

描述

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

类型

array

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

描述

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.

类型

array

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

类型

string

.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

描述

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.

类型

object

`.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces`

描述

`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".

类型

array

`.spec.redisProxy.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]`

类型

string

`.spec.redisProxy.config`

描述

`Config` provides additional configuration parameters for the Redis Proxy. Key-value pairs corresponding to proxy configuration options.

类型

object

`.spec.redisProxy.nodeSelector`

描述

`NodeSelector` constrains proxy pods to nodes with matching labels. See: <https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodeselector>

类型

object

.spec.redisProxy.resources

描述

Resources defines compute resource requirements for the proxy containers. Controls CPU and memory limits/requests.

类型

object

属性	类型	描述
		Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.
claims	array	This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. 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.redisProxy.resources.claims

描述

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.

类型

array

.spec.redisProxy.resources.claims[]

描述

ResourceClaim references one entry in PodSpec.ResourceClaims.

类型

object

必填

name

属性	类型	描述
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.redisProxy.resources.limits

描述

Limits describes the maximum amount of compute resources allowed. More info:

<https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/>

类型

object

.spec.redisProxy.resources.requests

描述

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/>

类型

object

.spec.redisProxy.tolerations

描述

Tolerations allow the proxy pods to be scheduled on nodes with matching taints. See:

<https://kubernetes.io/docs/concepts/scheduling-eviction/taint-and-toleration/>

类型

array

.spec.redisProxy.tolerations[]

描述

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

类型

object

属性	类型	描述
effect	string	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
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.
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.replicas

描述

Replicas defines the number of Redis nodes for different architectures. Configuration varies based on the selected architecture (cluster or sentinel).

类型

object

属性	类型	描述
cluster	object	Cluster holds replica settings specific to Redis Cluster architecture. Configures shard count and replica count per shard.
sentinel	object	Sentinel holds replica settings specific to Redis Sentinel architecture. Configures master and replica count.

.spec.replicas.cluster

描述

Cluster holds replica settings specific to Redis Cluster architecture. Configures shard count and replica count per shard.

类型

object

必填

shard

属性	类型	描述
<code>shard</code>	<code>integer</code>	Shard specifies the number of master nodes in the Redis Cluster. Each shard handles a subset of the hash slots (16384 total slots). Minimum value: 3
<code>shards</code>	<code>array</code>	Shards allows custom configuration of hash slot ranges for each shard. This configuration only works for newly created instances and cannot be updated after startup.
<code>slave</code>	<code>integer</code>	Slave specifies the number of replica nodes per master in the Redis Cluster. Each replica maintains a copy of its master's data for high availability. Range: 0-5

`.spec.replicas.cluster.shards`

描述

Shards allows custom configuration of hash slot ranges for each shard. This configuration only works for newly created instances and cannot be updated after startup.

类型

`array`

`.spec.replicas.cluster.shards[]`

类型

`object`

属性	类型	描述
<code>slots</code>	<code>string</code>	Slots is the slot range for the shard, eg: 0-1000,1002,1005-1100

.spec.replicas.sentinel

描述

Sentinel holds replica settings specific to Redis Sentinel architecture. Configures master and replica count.

类型

`object`

必填

`master`

属性	类型	描述
<code>master</code>	<code>integer</code>	Master specifies the number of Redis master nodes. For Sentinel architecture, only 1 master is supported.
<code>slave</code>	<code>integer</code>	Slave specifies the number of Redis replica nodes that replicate from the master. Range: 0-5

.spec.resources

描述

Resources defines the compute resource requirements for Redis containers. This includes CPU and memory limits/requests.

类型

object

属性	类型	描述
		Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.
claims	array	This is an alpha field and requires enabling the DynamicResourceAllocation feature gate. 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

描述

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.

类型

array

.spec.resources.claims[]

描述

ResourceClaim references one entry in PodSpec.ResourceClaims.

类型

object

必填

name

属性	类型	描述
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.resources.limits

描述

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

类型

object

.spec.resources.requests

描述

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/>

类型

object

.spec.restore

描述

Restore configures data restoration from a backup. Specifies the backup source and restoration parameters.

类型

object

属性	类型	描述
backupName	string	
image	string	
imagePullPolicy	string	PullPolicy describes a policy for if/when to pull a container image

.spec.securityContext

描述

SecurityContext defines security settings for Redis pods, including user/group IDs, filesystem permissions, and SELinux context.

类型

object

属性	类型	描述
appArmorProfile	object	appArmorProfile is the AppArmor options to use by the containers in this pod. Note that this field cannot be set when spec.os.name is windows.
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: 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---- 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

属性	类型	描述
		"Always". If not specified, "Always" is used. 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 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

属性	类型	描述
		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". "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. "MountOption" mounts all eligible Pod volumes with <code>-o context</code> 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. If not specified and SELinuxMount feature gate is enabled, "MountOption" is used. If not specified and SELinuxMount feature

属性	类型	描述
		gate is disabled, "MountOption" is used for ReadWriteOncePod volumes and "Recursive" for all other volumes. This field affects only Pods that have SELinux label set, either in PodSecurityContext or in SecurityContext of all containers. 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	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. 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.
seccompProfile	object	The seccomp options to use by the containers in this pod. Note that this field cannot be set when spec.os.name is windows.

属性	类型	描述
<code>supplementalGroups</code>	<code>array</code>	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.
<code>supplementalGroupsPolicy</code>	<code>string</code>	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.
<code>sysctls</code>	<code>array</code>	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail

属性	类型	描述
		to launch. Note that this field cannot be set when spec.os.name is windows.
windowsOptions	object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. 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 linux.

.spec.securityContext.appArmorProfile

描述

appArmorProfile is the AppArmor options to use by the containers in this pod. Note that this field cannot be set when spec.os.name is windows.

类型

object

必填

type

属性	类型	描述
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".

属性	类型	描述
<code>type</code>	<code>string</code>	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.

`.spec.securityContext.seLinuxOptions`

描述

The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. 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.

类型

`object`

属性	类型	描述
<code>level</code>	<code>string</code>	Level is SELinux level label that applies to the container.
<code>role</code>	<code>string</code>	Role is a SELinux role label that applies to the container.
<code>type</code>	<code>string</code>	Type is a SELinux type label that applies to the container.
<code>user</code>	<code>string</code>	User is a SELinux user label that applies to the container.

.spec.securityContext.seccompProfile

描述

The seccomp options to use by the containers in this pod. Note that this field cannot be set when spec.os.name is windows.

类型

object

必填

type

属性	类型	描述
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: 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.

.spec.securityContext.supplementalGroups

描述

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`.

类型

array

`.spec.securityContext.supplementalGroups[]`

类型

integer

`.spec.securityContext.sysctls`

描述

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`.

类型

array

`.spec.securityContext.sysctls[]`

描述

Sysctl defines a kernel parameter to be set

类型

object

必填

name

value

属性	类型	描述
name	string	Name of a property to set
value	string	Value of a property to set

.spec.securityContext.windowsOptions

描述

The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. 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 linux.

类型

object

属性	类型	描述
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 <code>GMSACredentialSpecName</code> field.
gmsaCredentialSpecName	string	GMSACredentialSpecName is the name of the GMSA credential spec to use.

属性	类型	描述
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.sentinel

描述

Sentinel configures Redis Sentinel for high availability. Only applicable when arch is set to "sentinel".

类型

object		
--------	--	--

属性	类型	描述
affinity	object	Affinity is a group of affinity scheduling rules.

属性	类型	描述
<code>customConfig</code>	<code>object</code>	Config the config for sentinel
<code>enableTLS</code>	<code>boolean</code>	EnableTLS enable TLS for redis
<code>exporter</code>	<code>object</code>	Exporter defines the specification for the sentinel exporter
<code>expose</code>	<code>object</code>	Expose
<code>externalTLSSecret</code>	<code>string</code>	ExternalTLSSecret the external TLS secret to use, if not provided, the operator will create one
<code>image</code>	<code>string</code>	Image the redis sentinel image
<code>imagePullPolicy</code>	<code>string</code>	ImagePullPolicy the image pull policy
<code>imagePullSecrets</code>	<code>array</code>	ImagePullSecrets the image pull secrets
<code>monitorConfig</code>	<code>object</code>	MonitorConfig configs for sentinel to monitor this replication, including: • <code>down-after-milliseconds</code> • <code>failover-timeout</code>

属性	类型	描述
		parallel-syncs
<code>nodeSelector</code>	<code>object</code>	
<code>passwordSecret</code>	<code>string</code>	PasswordSecret
<code>podAnnotations</code>	<code>object</code>	
<code>quorum</code>	<code>integer</code>	Quorum the number of Sentinels that need to agree about the fact the master is not reachable, in order to really mark the master as failing, and eventually start a failover procedure if possible. If not specified, the default value is the majority of the Sentinels.
<code>replicas</code>	<code>integer</code>	Replicas the number of sentinel replicas
<code>resources</code>	<code>object</code>	Resources the resources for sentinel
<code>securityContext</code>	<code>object</code>	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.
<code>sentinelReference</code>	<code>object</code>	SentinelReference the sentinel reference

属性	类型	描述
<code>serviceAnnotations</code>	<code>object</code>	
<code>tolerations</code>	<code>array</code>	

`.spec.sentinel.affinity`

描述

Affinity is a group of affinity scheduling rules.

类型

`object`

属性	类型	描述
<code>nodeAffinity</code>	<code>object</code>	Describes node affinity scheduling rules for the pod.
<code>podAffinity</code>	<code>object</code>	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
<code>podAntiAffinity</code>	<code>object</code>	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

`.spec.sentinel.affinity.nodeAffinity`

描述

Describes node affinity scheduling rules for the pod.

类型

object

属性	类型	描述
preferredDuringSchedulingIgnoredDuringExecution	array	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field. it may choose a node that violates one or more the expressions. The node that is most preferred is the one v the greatest sum of weights, i.e. for each node that meets all o scheduling requireme (resource request, requiredDuringSched affinity expressions, € compute a sum by iterating through the elements of this field adding "weight" to the sum if the node matc the corresponding matchExpressions; th node(s) with the high sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	object	If the affinity requiren specified by this field not met at scheduling

属性	类型	描述
		time, the pod will not scheduled onto the n If the affinity requiren specified by this field cease to be met at sc point during pod execution (e.g. due to update), the system r or may not try to eventually evict the p from its node.

.spec.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution

描述

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.

类型

array

.spec.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

描述

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).

类型

object

必填

preference

weight

属性	类型	描述
preference	object	A node selector term, associated with the corresponding weight.
weight	integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

.spec.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference

描述

A node selector term, associated with the corresponding weight.

类型

object

属性	类型	描述
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.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions`

描述

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

类型

array

`.spec.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[]`

描述

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

类型

object

必填

key

operator

属性	类型	描述
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.

属性	类型	描述
<code>values</code>	<code>array</code>	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.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values`

描述

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.

类型

`array`

`.spec.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchExpressions[].values[]`

类型

`string`

`.spec.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields`

描述

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

类型

array

.spec.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.
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.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values`

描述

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.

类型

`array`

`.spec.sentinel.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFields[].values[]`

类型

`string`

`.spec.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution`

描述

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 an update), the system may or may not try to eventually evict the pod from its node.

类型

`object`

必填

`nodeSelectorTerms`

属性	类型	描述
<code>nodeSelectorTerms</code>	<code>array</code>	Required. A list of node selector terms. The terms are ORed.

`.spec.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms`

描述

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

类型

`array`

`.spec.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]`

描述

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.

类型

`object`

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	A list of node selector requirements by node's labels.
<code>matchFields</code>	<code>array</code>	A list of node selector requirements by node's fields.

.spec.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions

描述

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

类型

array

.spec.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[]

描述

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

类型

object

必填

keyoperator

属性	类型	描述
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.

属性	类型	描述
<code>values</code>	<code>array</code>	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.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values`

描述

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.

类型

`array`

`.spec.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchExpressions[].values[]`

类型

`string`

`.spec.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFi`

elds

描述

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

类型

array

.spec.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.
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

属性	类型	描述
		element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

.spec.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values

描述

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.

类型

array

.spec.sentinel.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].matchFields[].values[]

类型

string

.spec.sentinel.affinity.podAffinity

描述

Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).

类型

object

属性	类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	<code>array</code>	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field. it may choose a node that violates one or more the expressions. The node that is most preferred is the one v the greatest sum of weights, i.e. for each node that meets all o scheduling requireme (resource request, requiredDuringSched affinity expressions, € compute a sum by iterating through the elements of this field adding "weight" to the sum if the node has r which matches the corresponding podAffinityTerm; the node(s) with the high sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	<code>array</code>	If the affinity requirem specified by this field not met at scheduling time, the pod will not

属性	类型	描述
		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 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 intersected, i.e. all terms must be satisfied.

`.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution`

描述

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.

类型

array

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

描述

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

类型

object

必填

podAffinityTerm weight

属性	类型	描述
podAffinityTerm	object	Required. A pod affinity term, associated with the corresponding weight.
weight	integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

描述

Required. A pod affinity term, associated with the corresponding weight.

类型

object

必填

topologyKey

属性	类型	描述
labelSelector	object	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
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 <code>labelSelector</code> as <code>key in (value)</code> 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 <code>labelSelector</code> as <code>key notin (value)</code> 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.

属性	类型	描述
		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 query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.
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.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

描述

A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.

类型

object

属性	类型	描述
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.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions

描述

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

类型

array

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values

描述

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.

类型

array

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]

类型

string

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels

描述

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.

类型

object

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

描述

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).

类型

array

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

类型

string

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys

描述

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).

类型

array

`.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]`

类型

string

`.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector`

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
<code>matchExpressions</code>	array	matchExpressions is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	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.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions`

描述

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

类型

array

`.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[]`

描述

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

类型

object

必填

key

operator

属性	类型	描述
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.

属性	类型	描述
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.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values

描述

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.

类型

array

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchExpressions[].values[]

类型

string

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaceSelector.matchLabels

描述

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.

类型

object

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

描述

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".

类型

array

.spec.sentinel.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces []

类型

string

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution

描述

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.

类型

array

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

描述

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

类型

object

必填

topologyKey

属性	类型	描述
labelSelector	object	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
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

属性	类型	描述
		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).
<code>mismatchLabelKeys</code>	<code>array</code>	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 <code>labelSelector</code> as <code>key notin (value)</code> 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).
<code>namespaceSelector</code>	<code>object</code>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

属性	类型	描述
<code>namespaces</code>	<code>array</code>	<code>namespaces</code> 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 <code>namespaceSelector</code> . null or empty <code>namespaces</code> list and null <code>namespaceSelector</code> means "this pod's namespace".
<code>topologyKey</code>	<code>string</code>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the <code>labelSelector</code> in the specified namespaces, where co-located is defined as running on a node whose value of the label with key <code>topologyKey</code> matches that of any node on which any of the selected pods is running. Empty <code>topologyKey</code> is not allowed.

`.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector`

描述

A label query over a set of resources, in this case pods. If it's null, this `PodAffinityTerm` matches with no Pods.

类型

`object`

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.

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

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions

描述

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

类型

`array`

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

描述

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

类型

`object`

必填

`key` `operator`

属性	类型	描述
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.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

描述

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.

类型

array

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

类型

`string`

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

描述

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.

类型

`object`

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

描述

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).

类型

`array`

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]

类型

`string`

`.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys`

描述

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).

类型

array

`.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]`

类型

string

`.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector`

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	<code>object</code>	<code>matchLabels</code> is a map of {key,value} pairs. A single {key,value} in the <code>matchLabels</code> map is equivalent to an element of <code>matchExpressions</code> , whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

`.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions`

描述

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

类型

`array`

`.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]`

描述

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

类型

`object`

必填

key

operator

属性	类型	描述
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.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

描述

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.

类型

array

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchEx

pressions[].values[]

类型

string

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

描述

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.

类型

object

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

描述

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".

类型

array

.spec.sentinel.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

类型

string

.spec.sentinel.affinity.podAntiAffinity

描述

Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

类型

object

属性	类型	描述
preferredDuringSchedulingIgnoredDuringExecution	array	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expression specified by this field. it may choose a node that violates one or more 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, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expression etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has a label which matches the

属性	类型	描述
		corresponding podAffinityTerm; the node(s) with the high sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	array	If the anti-affinity requirements specified in this field are not met at the time of pod scheduling time, the pod will not be scheduled on 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 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 intersected, i.e. all terms must be satisfied.

`.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution`

描述

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.

类型

array

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

描述

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

类型

object

必填

podAffinityTerm weight

属性	类型	描述
podAffinityTerm	object	Required. A pod affinity term, associated with the corresponding weight.
weight	integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm

描述

Required. A pod affinity term, associated with the corresponding weight.

类型

object

必填

topologyKey

属性	类型	描述
labelSelector	object	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
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).

属性	类型	描述
<code>mismatchLabelKeys</code>	<code>array</code>	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 <code>labelSelector</code> as <code>key notin (value)</code> 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).
<code>namespaceSelector</code>	<code>object</code>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.
<code>namespaces</code>	<code>array</code>	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".

属性	类型	描述
<code>topologyKey</code>	<code>string</code>	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.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector

描述

A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.

类型

`object`

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	matchExpressions is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	<code>object</code>	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.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions`

描述

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

类型

array

`.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[]`

描述

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

类型

object

必填

key

operator

属性	类型	描述
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.

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

`.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values`

描述

`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.

类型

`array`

`.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchExpressions[].values[]`

类型

`string`

`.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.labelSelector.matchLabels`

描述

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.

类型

object

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys

描述

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).

类型

array

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.matchLabelKeys[]

类型

string

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismat

chLabelKeys

描述

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).

类型

array

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.mismatchLabelKeys[]

类型

string

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespacesSelector

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
<code>matchExpressions</code>	<code>array</code>	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	<code>object</code>	<code>matchLabels</code> is a map of {key,value} pairs. A single {key,value} in the <code>matchLabels</code> map is equivalent to an element of <code>matchExpressions</code> , whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

`.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces.selector.matchExpressions`

描述

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

类型

`array`

`.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces.selector.matchExpressions[]`

描述

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

类型

`object`

必填

key

operator

属性	类型	描述
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.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces.selector.matchExpressions[].values

描述

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.

类型

array

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces.selector.matchExpressions[].values

namespaceSelector.matchExpressions[].values[]

类型

string

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

描述

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.

类型

object

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces

描述

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".

类型

array

.spec.sentinel.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm.namespaces[]

类型

string

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution

描述

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.

类型

array

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]

描述

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

类型

object

必填

topologyKey

属性	类型	描述
labelSelector	object	A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.
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 <code>labelSelector</code> as <code>key in (value)</code> 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 <code>labelSelector</code> as <code>key notin (value)</code> 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.

属性	类型	描述
		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 query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.
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.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector`

描述

A label query over a set of resources, in this case pods. If it's null, this PodAffinityTerm matches with no Pods.

类型

object

属性	类型	描述
<code>matchExpressions</code>	array	<code>matchExpressions</code> is a list of label selector requirements. The requirements are ANDed.
<code>matchLabels</code>	object	<code>matchLabels</code> 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.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions`

描述

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

类型

array

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values

描述

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.

类型

array

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchExpressions[].values[]

类型

string

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].labelSelector.matchLabels

描述

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.

类型

object

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys

描述

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).

类型

array

`.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].matchLabelKeys[]`

类型

string

`.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys`

描述

`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).

类型

array

`.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].mismatchLabelKeys[]`

类型

string

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector

描述

A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({}) matches all namespaces.

类型

object

属性	类型	描述
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.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions

描述

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

类型

array

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[]

描述

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

类型

object

必填

key operator

属性	类型	描述
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.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values

描述

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.

类型

array

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchExpressions[].values[]

类型

string

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaceSelector.matchLabels

描述

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.

类型

object

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces

描述

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".

类型

array

.spec.sentinel.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[].namespaces[]

类型

string

.spec.sentinel.customConfig

描述

Config the config for sentinel

类型

object

.spec.sentinel.exporter

描述

Exporter defines the specification for the sentinel exporter

类型

object

属性	类型	描述
<code>enabled</code>	<code>boolean</code>	
<code>image</code>	<code>string</code>	
<code>imagePullPolicy</code>	<code>string</code>	PullPolicy describes a policy for if/when to pull a container image

.spec.sentinel.expose

描述

Expose

类型

`object`

属性	类型	描述
<code>accessPort</code>	<code>integer</code>	AccessPort defines the lb access nodeport
<code>annotations</code>	<code>object</code>	The annnotations of the service which attached to services
<code>dataStorageNodePortMap</code>	<code>object</code>	NodePortMap defines the map of the nodeport for redis sentinel only Reversed for 3.14 backward compatibility

属性	类型	描述
<code>dataStorageNodePortSequence</code>	<code>string</code>	NodePortMap defines the map of the nodeport for redis nodes NodePortSequence defines the sequence of the nodeport for redis cluster only
<code>enableNodePort</code>	<code>boolean</code>	EnableNodePort defines if the nodeport is enabled
<code>image</code>	<code>string</code>	Image defines the image used to expose redis from annotations
<code>imagePullPolicy</code>	<code>string</code>	ImagePullPolicy defines the image pull policy
<code>ipFamilyPrefer</code>	<code>string</code>	IPFamily represents the IP Family (IPv4 or IPv6). This type is used to express the family of an IP expressed by a type (e.g. service.spec.ipFamilies).
<code>type</code>	<code>string</code>	ServiceType defines the type of the all related service

.spec.sentinel.expose.annotations

描述

The annotations of the service which attached to services

类型

object

.spec.sentinel.expose.dataStorageNodePortMap

描述

NodePortMap defines the map of the nodeport for redis sentinel only Reversed for 3.14 backward compatibility

类型

object

.spec.sentinel.imagePullSecrets

描述

ImagePullSecrets the image pull secrets

类型

array

.spec.sentinel.imagePullSecrets[]

描述

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

类型

object

属性	类型	描述
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

属性	类型	描述
		certainly wrong. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names ↗

.spec.sentinel.monitorConfig

描述

MonitorConfig configs for sentinel to monitor this replication, including: - down-after-milliseconds - failover-timeout - parallel-syncs

类型

object

.spec.sentinel.nodeSelector

类型

object

.spec.sentinel.podAnnotations

类型

object

.spec.sentinel.resources

描述

Resources the resources for sentinel

类型

object

属性	类型	描述
<code>claims</code>	<code>array</code>	Claims lists the names of resources, defined in <code>spec.resourceClaims</code>, that are used by this container. This is an alpha field and requires enabling the <code>DynamicResourceAllocation</code> feature gate. This field is immutable. It can only be set for containers.
<code>limits</code>	<code>object</code>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/ ↗
<code>requests</code>	<code>object</code>	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.sentinel.resources.claims`

描述

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.

类型

`array`

.spec.sentinel.resources.claims[]

描述

ResourceClaim references one entry in PodSpec.ResourceClaims.

类型

object

必填

name

属性	类型	描述
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.sentinel.resources.limits

描述

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

类型

object

.spec.sentinel.resources.requests

描述

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/>

类型

object

.spec.sentinel.securityContext

描述

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.

类型

object

属性	类型	描述
<code>appArmorProfile</code>	<code>object</code>	appArmorProfile is the AppArmor options to use by the containers in this pod. Note that this field cannot be set when spec.os.name is windows.
<code>fsGroup</code>	<code>integer</code>	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: 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)

属性	类型	描述
		3. The permission bits are OR'd with rw-rw- --- 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.
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.

属性	类型	描述
<code>runAsNonRoot</code>	<code>boolean</code>	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.
<code>runAsUser</code>	<code>integer</code>	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.
<code>seLinuxChangePolicy</code>	<code>string</code>	<code>seLinuxChangePolicy</code> 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". "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

属性	类型	描述
		Pods sharing the same volume on the same node. "MountOption" mounts all eligible Pod volumes with <code>-o context</code> 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 <code>spec.selinuxMount: true</code> in their CSIDriver instance. Other volumes are always re-labelled recursively. "MountOption" value is allowed only when SELinuxMount feature gate is enabled. 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. This field affects only Pods that have SELinux label set, either in PodSecurityContext or in SecurityContext of all containers. All Pods that use the same volume should use the same <code>seLinuxChangePolicy</code>, otherwise some pods can get stuck in ContainerCreating state. Note that this field cannot be set when <code>spec.os.name</code> is windows.

属性	类型	描述
<code>seLinuxOptions</code>	<code>object</code>	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. 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.
<code>seccompProfile</code>	<code>object</code>	The seccomp options to use by the containers in this pod. Note that this field cannot be set when spec.os.name is windows.

属性	类型	描述
<code>supplementalGroups</code>	<code>array</code>	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.
<code>supplementalGroupsPolicy</code>	<code>string</code>	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.
<code>sysctls</code>	<code>array</code>	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail

属性	类型	描述
		to launch. Note that this field cannot be set when spec.os.name is windows.
windowsOptions	object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. 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 linux.

.spec.sentinel.securityContext.appArmorProfile

描述

appArmorProfile is the AppArmor options to use by the containers in this pod. Note that this field cannot be set when spec.os.name is windows.

类型

object

必填

type

属性	类型	描述
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".

属性	类型	描述
<code>type</code>	<code>string</code>	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.

`.spec.sentinel.securityContext.seLinuxOptions`

描述

The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. 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.

类型

`object`

属性	类型	描述
<code>level</code>	<code>string</code>	Level is SELinux level label that applies to the container.
<code>role</code>	<code>string</code>	Role is a SELinux role label that applies to the container.
<code>type</code>	<code>string</code>	Type is a SELinux type label that applies to the container.
<code>user</code>	<code>string</code>	User is a SELinux user label that applies to the container.

.spec.sentinel.securityContext.seccompProfile

描述

The seccomp options to use by the containers in this pod. Note that this field cannot be set when spec.os.name is windows.

类型

object

必填

type

属性	类型	描述
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: 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.

.spec.sentinel.securityContext.supplementalGroups

描述

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.

类型

array

.spec.sentinel.securityContext.supplementalGroups[]

类型

integer

.spec.sentinel.securityContext.sysctls

描述

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.

类型

array

.spec.sentinel.securityContext.sysctls[]

描述

Sysctl defines a kernel parameter to be set

类型

object

必填

name

value

属性	类型	描述
name	string	Name of a property to set
value	string	Value of a property to set

.spec.sentinel.securityContext.windowsOptions

描述

The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. 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 linux.

类型

object

属性	类型	描述
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 <code>GMSACredentialSpecName</code> field.
gmsaCredentialSpecName	string	GMSACredentialSpecName is the name of the GMSA credential spec to use.

属性	类型	描述
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.sentinel.sentinelReference

描述

SentinelReference the sentinel reference

类型

object

属性	类型	描述
auth	object	Auth the sentinel auth

属性	类型	描述
<code>nodes</code>	<code>array</code>	Addresses the sentinel addresses

`.spec.sentinel.sentinelReference.auth`

描述

Auth the sentinel auth

类型

`object`

属性	类型	描述
<code>passwordSecret</code>	<code>string</code>	PasswordSecret the password secret for redis
<code>tlsSecret</code>	<code>string</code>	TLSSecret the tls secret
<code>username</code>	<code>string</code>	Username the username for redis

`.spec.sentinel.sentinelReference.nodes`

描述

Addresses the sentinel addresses

类型

`array`

`.spec.sentinel.sentinelReference.nodes[]`

类型

object

属性	类型	描述
flags	string	Flags
ip	string	IP the sentinel node ip
port	integer	Port the sentinel node port

.spec.sentinel.serviceAnnotations

类型

object

.spec.sentinel.tolerations

类型

array

.spec.sentinel.tolerations[]

描述

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

类型

object

属性	类型	描述
effect	string	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
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.
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.sentinelCustomConfig

描述

SentinelCustomConfig defines custom Redis Sentinel configuration settings. Deprecated: Use sentinel.customConfig instead for improved organization.

类型

object

.spec.tolerations

描述

Tolerations allow Redis pods to be scheduled on nodes with matching taints. See: <https://kubernetes.io/docs/concepts/scheduling-eviction/taint-and-toleration/>

类型

array

.spec.tolerations[]

描述

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

类型

object

属性	类型	描述
effect	string	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.

属性	类型	描述
<code>key</code>	<code>string</code>	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.
<code>operator</code>	<code>string</code>	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.
<code>tolerationSeconds</code>	<code>integer</code>	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.
<code>value</code>	<code>string</code>	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.upgradeOption

描述

UpgradeOption defines the upgrade strategy and automation settings. Controls how and when Redis instances are upgraded.

类型

object

属性	类型	描述
autoUpgrade	boolean	AutoUpgrade controls whether the operator should automatically upgrade the instance when a new version becomes available.
crVersion	string	CRVersion specifies the target version for the Redis instance to upgrade to.

.status

描述

RedisStatus defines the observed state of a Redis instance.

类型

object

属性	类型	描述
clusterNodes	array	ClusterNodes contains detailed information about each Redis node in the cluster. Includes role, address, and health status.
detailedStatusRef	object	DetailedStatusRef references a resource with extended status information. Used for complex status data that doesn't fit in the main status.

属性	类型	描述
<code>lastShardCount</code>	<code>integer</code>	LastShardCount records the previous number of shards in the Redis Cluster. Used to detect and handle scaling operations.
<code>lastVersion</code>	<code>string</code>	LastVersion records the previous version of the Redis instance. Used to detect and handle version upgrades.
<code>matchLabels</code>	<code>object</code>	MatchLabels contains the selector labels used to identify Redis pods.
<code>message</code>	<code>string</code>	Message provides additional details about the current status. Contains error information or in-progress operation details.
<code>passwordSecretName</code>	<code>string</code>	PasswordSecretName is the name of the Kubernetes Secret containing the Redis password. This is either user-specified or automatically generated.

属性	类型	描述
<code>phase</code>	<code>string</code>	Phase indicates the current operational state of the Redis instance. Values: Initializing - Resources are being created or reconciled Ready - All resources are ready and Redis is operational Error - An error occurred during resource initialization or operation Rebalancing - Cluster slots are being redistributed Paused - Reconciliation is paused
<code>proxyMatchLabels</code>	<code>object</code>	ProxyMatchLabels contains the selector labels used to identify Redis proxy pods.
<code>proxyServiceName</code>	<code>string</code>	ProxyServiceName is the name of the Kubernetes Service created for Redis Proxy.
<code>restored</code>	<code>boolean</code>	Restored indicates whether the instance has been successfully restored from a backup. Set to true once a restore operation completes successfully.
<code>serviceName</code>	<code>string</code>	ServiceName is the name of the main Kubernetes Service created for Redis access.
<code>upgradeStatus</code>	<code>object</code>	UpgradeStatus contains detailed information about the ongoing or last upgrade operation.

.status.clusterNodes

描述

ClusterNodes contains detailed information about each Redis node in the cluster. Includes role, address, and health status.

类型

array

.status.clusterNodes[]

描述

RedisNode represent a RedisCluster Node

类型

object

必填

ip nodeName podName port role statefulSet

属性	类型	描述
id	string	ID is the redis cluster node id, not runid
ip	string	IP is the ip of the node. if access announce is enabled, it will be the access ip
masterRef	string	MasterRef is the master node id of this node
nodeName	string	nodeName is the node name of the node where holds the pod

属性	类型	描述
<code>podName</code>	<code>string</code>	PodName current pod name
<code>port</code>	<code>string</code>	Port is the port of the node. if access announce is enabled, it will be the access port
<code>role</code>	<code>string</code>	Role is the role of the node, master or slave
<code>slots</code>	<code>array</code>	Slots is the slot range for the shard, eg: 0-1000,1002,1005-1100
<code>statefulSet</code>	<code>string</code>	StatefulSet is the statefulset name of this pod

.status.clusterNodes[].slots

描述

Slots is the slot range for the shard, eg: 0-1000,1002,1005-1100

类型

`array`

.status.clusterNodes[].slots[]

类型

`string`

.status.detailedStatusRef

描述

DetailedStatusRef references a resource with extended status information. Used for complex status data that doesn't fit in the main status.

类型

object

属性	类型	描述
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 <code>desiredState.manifest.containers[2]</code>. For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (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

属性	类型	描述
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 ↗
uid	string	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids ↗

.status.matchLabels

描述

MatchLabels contains the selector labels used to identify Redis pods.

类型

object

.status.proxyMatchLabels

描述

ProxyMatchLabels contains the selector labels used to identify Redis proxy pods.

类型

object

.status.upgradeStatus

描述

UpgradeStatus contains detailed information about the ongoing or last upgrade operation.

类型

object

属性	类型	描述
crVersion	string	CRVersion indicates the target version for the upgrade operation.
message	string	Message provides additional information about the current upgrade status.

API 端点

可用的 API 端点如下：

- `/apis/middleware.alauda.io/v1/namespaces/{namespace}/redis`
 - `DELETE` : delete collection of Redis
 - `GET` : list objects of kind Redis
 - `POST` : create a new Redis
- `/apis/middleware.alauda.io/v1/namespaces/{namespace}/redis/{name}`
 - `DELETE` : delete the specified Redis
 - `GET` : read the specified Redis
 - `PATCH` : partially update the specified Redis
 - `PUT` : replace the specified Redis

- `/apis/middleware.alauda.io/v1/namespaces/{namespace}/redis/{name}/status`
 - `GET` : read status of the specified Redis
 - `PATCH` : partially update status of the specified Redis
 - `PUT` : replace status of the specified Redis

/apis/middleware.alauda.io/v1/namespaces/{namespace}/redis

HTTP 方法

DELETE

描述

delete collection of Redis

HTTP 响应

HTTP 状态码	响应体
200 - OK	<code>Status</code> schema
401 - Unauthorized	Empty

HTTP 方法

GET

描述

list objects of kind Redis

HTTP 响应

HTTP 状态码	响应体
200 - OK	<code>RedisList</code> schema
401 - Unauthorized	Empty

HTTP 方法

POST

描述

create a new Redis

查询参数

参数	类型	描述
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.

请求体参数

参数	类型	描述
body	Redis schema	application/json formatted

HTTP 响应

HTTP 状态码	响应体
200 - OK	Redis schema
201 - Created	Redis schema
202 - Accepted	Redis schema
401 - Unauthorized	Empty

/apis/middleware.alauda.io/v1/namespaces/{namespace}/redis/{name}

HTTP 方法

DELETE

描述

delete the specified Redis

查询参数

参数	类型	描述
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 响应

HTTP 状态码	响应体
200 - OK	Status  schema
202 - Accepted	Status  schema
401 - Unauthorized	Empty

HTTP 方法

GET

描述

read the specified Redis

HTTP 响应

HTTP 状态码	响应体
200 - OK	<code>Redis</code> schema
401 - Unauthorized	Empty

HTTP 方法

PATCH

描述

partially update the specified Redis

查询参数

参数	类型	描述
<code>dryRun</code>	<code>string</code>	When present, indicates that modifications should not be persisted. An invalid or unrecognized <code>dryRun</code> 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
<code>fieldValidation</code>	<code>string</code>	<code>fieldValidation</code> 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.

参数	类型	描述
		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 响应

HTTP 状态码	响应体
200 - OK	<code>Redis</code> schema
401 - Unauthorized	Empty

HTTP 方法

PUT

描述

replace the specified Redis

查询参数

参数	类型	描述
<code>dryRun</code>	<code>string</code>	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
<code>fieldValidation</code>	<code>string</code>	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,

参数	类型	描述
		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.

请求体参数

参数	类型	描述
<code>body</code>	<code>Redis</code> schema	<code>application/json</code> formatted

HTTP 响应

HTTP 状态码	响应体
200 - OK	<code>Redis</code> schema
201 - Created	<code>Redis</code> schema
401 - Unauthorized	Empty

/apis/middleware.alauda.io/v1/namespaces/{namespace}/redis/{name}/status

HTTP 方法

GET

描述

read status of the specified Redis

HTTP 响应

HTTP 状态码	响应体
200 - OK	<code>Redis</code> schema
401 - Unauthorized	Empty

HTTP 方法

PATCH

描述

partially update status of the specified Redis

查询参数

参数	类型	描述
<code>dryRun</code>	<code>string</code>	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
<code>fieldValidation</code>	<code>string</code>	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.

参数	类型	描述
<code>dryRun</code>	<code>string</code>	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

参数	类型	描述
<code>fieldValidation</code>	<code>string</code>	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.

请求体参数

参数	类型	描述
<code>body</code>	<code>Redis</code> schema	<code>application/json</code> formatted

HTTP 响应

HTTP 状态码	响应体
200 - OK	<code>Redis</code> schema
201 - Created	<code>Redis</code> schema
401 - Unauthorized	Empty



Redis 用户

描述

RedisUser is the Schema for the redisusers API. Represents a Redis user account with associated ACL permissions.

类型

object

规格

属性	类型	描述
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:

属性	类型	描述
		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	RedisUserSpec defines the desired state of RedisUser
status	object	RedisUserStatus defines the observed state of RedisUser

.spec

描述

RedisUserSpec defines the desired state of RedisUser

类型

object

必填

redisName username

属性	类型	描述
accountType	string	AccountType defines the user account category. Options: system (privileged system accounts), custom (general user accounts), default (default user configuration)

属性	类型	描述
<code>aclRules</code>	<code>string</code>	AclRules defines the Redis ACL rules as a string. Format follows Redis ACL syntax for permissions and command restrictions. See: https://redis.io/docs/management/security/acl/
<code>arch</code>	<code>string</code>	Arch specifies the Redis architecture type for this user. Must match the architecture of the referenced Redis instance.
<code>passwordSecrets</code>	<code>array</code>	PasswordSecrets is a list of Kubernetes Secret names containing user passwords. Each secret must have a 'password' key containing the Redis password. Multiple secrets allow for credential rotation.
<code>redisName</code>	<code>string</code>	RedisName specifies the target Redis instance name. Must reference an existing Redis resource in the same namespace.
<code>username</code>	<code>string</code>	Username specifies the Redis user account name. Required for user identification during authentication.

.spec.passwordSecrets

描述

PasswordSecrets is a list of Kubernetes Secret names containing user passwords. Each secret must have a 'password' key containing the Redis password. Multiple secrets allow for credential rotation.

类型

array

.spec.passwordSecrets[]

类型

string

.status

描述

RedisUserStatus defines the observed state of RedisUser

类型

object

属性	类型	描述
Phase	string	Phase indicates the current state of the RedisUser resource. Values: Fail, Success, Pending
aclRules	string	AclRules reflects the actual ACL rules applied to Redis.
lastUpdateSuccess	string	LastUpdatedSuccess records the timestamp of the last successful configuration update. Format: RFC3339 timestamp string.

属性	类型	描述
<code>message</code>	<code>string</code>	Message provides additional details about the current status. Contains error information or operation details.

API 端点

可用的 API 端点如下：

- `/apis/redis.middleware.alauda.io/v1/namespaces/{namespace}/redisusers`
 - `DELETE` : delete collection of RedisUser
 - `GET` : list objects of kind RedisUser
 - `POST` : create a new RedisUser
- `/apis/redis.middleware.alauda.io/v1/namespaces/{namespace}/redisusers/{name}`
 - `DELETE` : delete the specified RedisUser
 - `GET` : read the specified RedisUser
 - `PATCH` : partially update the specified RedisUser
 - `PUT` : replace the specified RedisUser
- `/apis/redis.middleware.alauda.io/v1/namespaces/{namespace}/redisusers/{name}/status`
 - `GET` : read status of the specified RedisUser
 - `PATCH` : partially update status of the specified RedisUser
 - `PUT` : replace status of the specified RedisUser

/apis/redis.middleware.alauda.io/v1/namespaces/{namespace}/redisusers

HTTP 方法

[DELETE](#)

描述

delete collection of RedisUser

HTTP 响应

HTTP 状态码	响应体
200 - OK	Status schema
401 - Unauthorized	Empty

HTTP 方法

[GET](#)

描述

list objects of kind RedisUser

HTTP 响应

HTTP 状态码	响应体
200 - OK	RedisUserList schema
401 - Unauthorized	Empty

HTTP 方法

[POST](#)

描述

create a new RedisUser

查询参数

参数	类型	描述
<code>dryRun</code>	<code>string</code>	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
<code>fieldValidation</code>	<code>string</code>	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.

请求体参数

参数	类型	描述
<code>body</code>	<code>RedisUser</code> schema	<code>application/json</code> formatted

HTTP 响应

HTTP 状态码	响应体
200 - OK	<code>RedisUser</code> schema
201 - Created	<code>RedisUser</code> schema

HTTP 状态码	响应体
202 - Accepted	RedisUser schema
401 - Unauthorized	Empty

/apis/redis.middleware.alauda.io/v1/namespaces/{namespace}/redisusers/{name}

HTTP 方法

DELETE

描述

delete the specified RedisUser

查询参数

参数	类型	描述
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 响应

HTTP 状态码	响应体
200 - OK	Status  schema
202 - Accepted	Status  schema
401 - Unauthorized	Empty

HTTP 方法

GET

描述

read the specified RedisUser

HTTP 响应

HTTP 状态码	响应体
200 - OK	<code>RedisUser</code> schema
401 - Unauthorized	Empty

HTTP 方法

`PATCH`

描述

partially update the specified RedisUser

查询参数

参数	类型	描述
<code>dryRun</code>	<code>string</code>	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
<code>fieldValidation</code>	<code>string</code>	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

参数	类型	描述
		will contain all unknown and duplicate fields encountered.

HTTP 响应

HTTP 状态码	响应体
200 - OK	<code>RedisUser</code> schema
401 - Unauthorized	Empty

HTTP 方法

PUT

描述

replace the specified RedisUser

查询参数

参数	类型	描述
<code>dryRun</code>	<code>string</code>	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
<code>fieldValidation</code>	<code>string</code>	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

参数	类型	描述
		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.

请求体参数

参数	类型	描述
body	RedisUser schema	application/json formatted

HTTP 响应

HTTP 状态码	响应体
200 - OK	RedisUser schema
201 - Created	RedisUser schema
401 - Unauthorized	Empty

/apis/redis.middleware.alauda.io/v1/namespaces/{namespace}/redisusers/{name}/status

HTTP 方法

GET

描述

read status of the specified RedisUser

HTTP 响应

HTTP 状态码	响应体
200 - OK	RedisUser schema

HTTP 状态码	响应体
401 - Unauthorized	Empty

HTTP 方法

PATCH

描述

partially update status of the specified RedisUser

查询参数

参数	类型	描述
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 响应

HTTP 状态码	响应体
200 - OK	<code>RedisUser</code> schema
401 - Unauthorized	Empty

HTTP 方法

PUT

描述

replace status of the specified RedisUser

查询参数

参数	类型	描述
<code>dryRun</code>	<code>string</code>	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
<code>fieldValidation</code>	<code>string</code>	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.

请求体参数

参数	类型	描述
body	RedisUser schema	application/json formatted

HTTP 响应

HTTP 状态码	响应体
200 - OK	RedisUser schema
201 - Created	RedisUser schema
401 - Unauthorized	Empty