Token APIs

AccessTokenInfo [auth.alauda.io/v1]

AccessTokenInfo [auth.alauda.io/v1]

/apis/auth.alauda.io/v1/namespaces/{namespace}/accesstokeninfoes

公共参数

- `namespace` (*in path*): `string` required
  object name and auth scope, such as for teams and projects
- `pretty` (*in query*): `string`
  If 'true', then the output is pretty printed. Defaults to 'false' unless the user-agent indicates a browser or command-line HTTP tool (curl and wget).

`get`

list objects of kind AccessTokenInfo

参数

- `allowWatchBookmarks` (*in query*): `boolean`
  allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

- `continue` (*in query*): `string`
  The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".
  This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.
- `fieldSelector` (*in query*): `string`
  A selector to restrict the list of returned objects by their fields. Defaults to everything.
- `labelSelector` (*in query*): `string`
  A selector to restrict the list of returned objects by their labels. Defaults to everything.
- `limit` (*in query*): `integer`
  limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.
  The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.

- `resourceVersion` (*in query*): `string`
  resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions ↗ for details.
  Defaults to unset
- `resourceVersionMatch` (*in query*): `string`
  resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions ↗ for details.
  Defaults to unset
- `sendInitialEvents` (*in query*): `boolean`
  `sendInitialEvents=true` may be set together with `watch=true`. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with `"k8s.io/initial-events-end": "true"` annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.
  When `sendInitialEvents` option is set, we require `resourceVersionMatch` option to also be set. The semantic of the watch request is as following: - `resourceVersionMatch` = NotOlderThan is interpreted as "data at least as new as the provided `resourceVersion`" and the bookmark event is send when the state is synced to a `resourceVersion` at least as fresh as the one provided by the ListOptions. If `resourceVersion` is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed.
    - `resourceVersionMatch` set to any other value or unset Invalid error is returned.

  Defaults to true if `resourceVersion=""` or `resourceVersion="0"` (for backward compatibility reasons) and to false otherwise.
- `timeoutSeconds` (*in query*): `integer`
  Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
- `watch` (*in query*): `boolean`

Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

响应

- `200` AccessTokenInfoList: OK
- `401` : Unauthorized

`delete`

delete collection of AccessTokenInfo

参数

- `allowWatchBookmarks` (*in query*): `boolean`
  allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
- `continue` (*in query*): `string`
  The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".
  This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.
- `fieldSelector` (*in query*): `string`

A selector to restrict the list of returned objects by their fields. Defaults to everything.

- `labelSelector` (*in query*): `string`
  A selector to restrict the list of returned objects by their labels. Defaults to everything.
- `limit` (*in query*): `integer`
  limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.
  The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.
- `resourceVersion` (*in query*): `string`
  resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions ↗ for details.
  Defaults to unset
- `resourceVersionMatch` (*in query*): `string`
  resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions ↗ for details.
  Defaults to unset
- `sendInitialEvents` (*in query*): `boolean`
  `sendInitialEvents=true` may be set together with `watch=true`. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent.

The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with `"k8s.io/initial-events-end": "true"` annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.
When `sendInitialEvents` option is set, we require `resourceVersionMatch` option to also be set. The semantic of the watch request is as following: - `resourceVersionMatch` = NotOlderThan is interpreted as "data at least as new as the provided `resourceVersion`" and the bookmark event is send when the state is synced to a `resourceVersion` at least as fresh as the one provided by the ListOptions. If `resourceVersion` is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed.

- `resourceVersionMatch` set to any other value or unset Invalid error is returned.

Defaults to true if `resourceVersion=""` or `resourceVersion="0"` (for backward compatibility reasons) and to false otherwise.

- `timeoutSeconds` (*in query*): `integer`
  Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
- `watch` (*in query*): `boolean`
  Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

响应

- `200` Status: OK
- `401` : Unauthorized

/apis/auth.alauda.io/v1/namespaces/{namespace}/accesstokeninfoes/{name}

公共参数

- `name` (*in path*): `string` required
  name of the AccessTokenInfo
- `namespace` (*in path*): `string` required
  object name and auth scope, such as for teams and projects
- `pretty` (*in query*): `string`
  If 'true', then the output is pretty printed. Defaults to 'false' unless the user-agent indicates a browser or command-line HTTP tool (curl and wget).

`get`

read the specified AccessTokenInfo

参数

- `resourceVersion` (*in query*): `string`
  resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions ↗ for details.
  Defaults to unset

响应

- `200` AccessTokenInfo: OK
- `401` : Unauthorized

`delete`

delete an AccessTokenInfo

参数

- `dryRun` (*in query*): `string`
  When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
- `gracePeriodSeconds` (*in query*): `integer`

The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

- `ignoreStoreReadErrorWithClusterBreakingPotential` (*in query*): `boolean`
  if set to true, it will trigger an unsafe deletion of the resource in case the normal deletion flow fails with a corrupt object error. A resource is considered corrupt if it can not be retrieved from the underlying storage successfully because of a) its data can not be transformed e.g. decryption failure, or b) it fails to decode into an object. NOTE: unsafe deletion ignores finalizer constraints, skips precondition checks, and removes the object from the storage. WARNING: This may potentially break the cluster if the workload associated with the resource being unsafe-deleted relies on normal deletion flow. Use only if you REALLY know what you are doing. The default value is false, and the user must opt in to enable it
- `orphanDependents` (*in query*): `boolean`
  Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
- `propagationPolicy` (*in query*): `string`
  Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

请求体

DeleteOptions

响应

- `200` Status: OK
- `202` Status: Accepted
- `401` : Unauthorized

AccessTokenInfoList

AccessTokenInfoList is a list of AccessTokenInfo

- `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 ↗
- `items`: `[]AccessTokenInfo`
  List of accesstokeninfoes. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md ↗
- `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`: ListMeta
  ListMeta describes metadata that synthetic resources must have, including lists and various status objects. A resource may have only one of {ObjectMeta, ListMeta}.

AccessTokenInfo

AccessTokenInfo is the Schema for the accesstokeninfoes API

- `apiVersion`: `string`
  APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources ↗
- `kind`: `string`
  Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In

CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds ↗

- `metadata` : ObjectMeta
  ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
- `spec` : `object`
  AccessTokenInfoSpec defines the desired state of AccessTokenInfo
- `status` : `object`
  AccessTokenInfoStatus defines the observed state of AccessTokenInfo

ListMeta

ListMeta describes metadata that synthetic resources must have, including lists and various status objects. A resource may have only one of {ObjectMeta, ListMeta}.

- `continue` : `string`
  continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
- `remainingItemCount` : `integer`
  remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is *estimating* the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
- `resourceVersion` : `string`

String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency ↗

- `selfLink` : `string`
  Deprecated: selfLink is a legacy read-only field that is no longer populated by the system.

Status

Status is a return value for calls that don't return other objects.

- `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 ↗
- `code` : `integer`
  Suggested HTTP return code for this status, 0 if not set.
- `details` : StatusDetails
  StatusDetails is a set of additional properties that MAY be set by the server to provide additional information about a response. The Reason field of a Status object defines what attributes will be set. Clients must ignore fields that do not match the defined type of each attribute, and should assume that any attribute may be empty, invalid, or under defined.
- `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 ↗
- `message` : `string`
  A human-readable description of the status of this operation.
- `metadata` : ListMeta
  ListMeta describes metadata that synthetic resources must have, including lists and various status objects. A resource may have only one of {ObjectMeta, ListMeta}.

- `reason` : `string`
  A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
- `status` : `string`
  Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status ↗

StatusDetails

StatusDetails is a set of additional properties that MAY be set by the server to provide additional information about a response. The Reason field of a Status object defines what attributes will be set. Clients must ignore fields that do not match the defined type of each attribute, and should assume that any attribute may be empty, invalid, or under defined.

- `causes` : `[]StatusCause`
  The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
- `group` : `string`
  The group attribute of the resource associated with the status StatusReason.
- `kind` : `string`
  The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds ↗
- `name` : `string`
  The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
- `retryAfterSeconds` : `integer`
  If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action - for those errors this field may indicate how long to wait before taking the alternate action.
- `uid` : `string`

UID of the resource. (when there is a single resource which can be described). More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names#uids ↗

StatusCause

StatusCause provides more information about an api.Status failure, including cases when multiple errors are encountered.

- `field` : `string`
  The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional.
  Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
- `message` : `string`
  A human-readable description of the cause of the error. This field may be presented as-is to a reader.
- `reason` : `string`
  A machine-readable description of the cause of the error. If this value is empty there is no information available.