



# **Unified Manager REST APIs**

## Active IQ Unified Manager 9.9

NetApp  
January 31, 2025

This PDF was generated from <https://docs.netapp.com/us-en/active-iq-unified-manager-99/api-automation/concept-data-center-apis.html> on January 31, 2025. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

Unified Manager REST APIs .....	1
Managing storage objects in a data center .....	1
Accessing ONTAP APIs through proxy access .....	6
Performing administrative tasks .....	8
Managing users .....	10
Viewing performance metrics .....	11
Viewing Jobs .....	16
Viewing events and system details .....	17
Managing workloads .....	18

# Unified Manager REST APIs

The REST APIs for Active IQ Unified Manager are listed in this section, based on their categories.

You can view the online documentation page from your Unified Manager instance that includes the details of every REST API call. This document does not repeat the details of the online documentation. Each API call listed or described in this document includes only the information you need to locate the call on the documentation page. After locating a specific API call, you can review the complete details of that call, including the input parameters, output formats, HTTP status codes, and request processing type.

The following information is included for each API call within a workflow to help locate the call on the documentation page:

- Category

The API calls are organized on the documentation page into functionally-related areas or categories. To locate a specific API call, scroll down to the bottom of the page, and then click the applicable API category.

- HTTP verb (call)

The HTTP verb identifies the action performed on a resource. Each API call is executed through a single HTTP verb.

- Path

The path determines the specific resource which the action uses to as a part of performing a call. The path string is appended to the core URL to form the complete URL identifying the resource.

## Managing storage objects in a data center

The REST APIs under the `datacenter` category enable you to manage the storage objects in your data center, such as clusters, nodes, aggregates, storage VMs, volumes, LUNs, file shares, and namespaces. These APIs are available for querying the configuration of the objects, while some of them enable you to perform the operations of adding, deleting, or modifying these objects.

Most of these APIs are GET calls that provide cross-cluster aggregation with filtering, sorting and pagination support. On running these APIs, they return data from the database. Therefore, the newly created-objects need to be discovered by the next acquisition cycle for appearing in the response.

If you want to query the details of a specific object, you need to enter the unique ID of that object to view its details. For example,

```
curl -X GET "https://<hostname>/api/datacenter/cluster/clusters/4c6bf721-2e3f-11e9-a3e2-00a0985badbb" -H "accept: application/json" -H "Authorization: Basic <Base64EncodedCredentials>"
```



The CURL commands, examples, requests, and responses to the APIs, are available on your Swagger API interface. You can filter and order the results by specific parameters as indicated on Swagger. These APIs enable you to filter the results for specific storage objects, such as cluster, volume, or storage VM.

HTTP verb	Path	Description
GET	/datacenter/cluster/clusters /datacenter/cluster/clusters/{key}	You can use this method to view the details of the ONTAP clusters across the data center. The API returns information, such as the IPv4 or IPv6 address of the cluster, information about the node, such as node health, performance capacity, and High Availability (HA) pair, and indicates whether the cluster is All SAN Array.
GET	/datacenter/cluster/nodes /datacenter/cluster/nodes/{key}	You can use this method to view the details of the nodes in the data center. You can view information about the cluster, node health, performance capacity, and High Availability (HA) pair for the node.
GET	/datacenter/protocols/cifs/shares /datacenter/protocols/cifs/shares/{key}	You can use this method to view the details of the CIFS shares in the data center. Apart from cluster, SVM, and volume details, information about Access Control List (ACL) is also returned.
GET	/datacenter/protocols/nfs/export-policies /datacenter/protocols/nfs/export-policies/{key}	You can use this method to view the details of the export policies for the supported NFS services.  You can query the export policies for a cluster or storage VM and reuse the export policy key for provisioning NFS file shares. For more information about the assigning and reusing export policies on workloads, see “Provisioning CIFS and NFS file shares”.

HTTP verb	Path	Description
GET	<p>/datacenter/storage/aggregates</p> <p>/datacenter/storage/aggregates/{key}</p>	<p>You can use this method to view the collection of aggregates in the data center or a specific aggregate for provisioning workloads on them or monitoring. Information, such as cluster and node details, performance capacity used, available and used space, and storage efficiency is returned.</p>
GET	<p>/datacenter/storage/luns</p> <p>/datacenter/storage/luns/{key}</p>	<p>You can use this method to view the collection of LUNs in the entire data center. You can view information about the LUN, such as cluster and SVM details, QoS policies, and igroups.</p>
GET	<p>/datacenter/storage/qos/policies</p> <p>/datacenter/storage/qos/policies/{key}</p>	<p>You can use this method to view the details of all the QoS policies applicable for the storage objects in the data center. Information, such as the cluster and SVM details, the fixed or adaptive policy details, and number of objects applicable for that policy is returned.</p>
GET	<p>/datacenter/storage/qtrees</p> <p>/datacenter/storage/qtrees/{key}</p>	<p>You can use this method to view the qtree details across the data center for all FlexVol volumes or FlexGroup volumes. Information, such as the cluster and SVM details, FlexVol volume, and export policy are returned.</p>
GET	<p>/datacenter/storage/volumes</p> <p>/datacenter/storage/volumes/{key}</p>	<p>You can use this method to view the collection of volumes in the data center. Information about the volumes, such as SVM and cluster details, QoS and export policies, whether the volume is of type read-write, data-protection, or load-sharing, is returned.</p> <p>For FlexVol and FlexClone volumes, you can view the information about the respective aggregates. For a FlexGroup volume, the query returns the list of constituent aggregates.</p>

HTTP verb	Path	Description
GET	/datacenter/protocols/san/igroups	
POST		
DELETE	/datacenter/protocols/san/igroups/{key}	
PATCH		<p>You can assign initiator groups (igroups) authorized to access particular LUN targets. If there is an existing igroup, you can assign it. You can also create igroups and assign them to the LUNs.</p>
		<p>You can use these methods to query, create, delete, and modify igroups respectively.</p> <p>Points to note:</p> <ul style="list-style-type: none"> <li>• POST: While creating an igroup, you can designate the storage VM on which you want to assign access.</li> <li>• DELETE: You need to provide the igroup key as an input parameter to delete a particular ingroup. If you have already assigned an ingroup to a LUN, you cannot delete that ingroup.</li> <li>• PATCH: You need to provide the ingroup key as an input parameter to modify a particular ingroup. You must also enter the property that you want to update, along with its value.</li> </ul>

HTTP verb	Path	Description
GET	/datacenter/svm/svms	You can use these methods to view, create, delete, and modify Storage Virtual Machines (storage VMs).
POST	/datacenter/svm/svms/{key}	
DELETE		Points to note:
PATCH		<ul style="list-style-type: none"> <li>POST: You must enter the storage VM object that you want to create as an input parameter. You can create a custom storage VM, and then assign required properties to it.</li> </ul>
		<p>If you have enabled SLO-based workload provisioning on your environment, while creating the storage VM, ensure that it supports all of the protocols required for provisioning LUNs and file shares on them, for example, CIFS or SMB, NFS, FCP, and iSCSI. The provisioning workflows might fail if the storage VM does not support the required services. It is recommended that the services for the respective types of workloads are also enabled on the storage VM.</p>
		<p> • DELETE: You need to provide the storage VM key to delete a particular storage VM.</p>

# Accessing ONTAP APIs through proxy access

The gateway APIs provide you with the advantage of using the Active IQ Unified Manager credentials to run ONTAP REST APIs and managing storage objects. These APIs are available when the API Gateway feature is enabled from the Unified Manager web UI.

Unified Manager REST APIs support only a select set of actions to be performed on the Unified Manager datasources, that is ONTAP clusters. You can avail the other features through ONTAP APIs. The gateway APIs allow Unified Manager to be a pass-through interface for tunneling all API requests to be performed on ONTAP clusters, without logging in to each data center cluster individually. It performs as a single point of management for running the APIs across the ONTAP clusters managed by your Unified Manager instance. The API Gateway feature allows Unified Manager to be a single control plane from which you can manage multiple ONTAP clusters, without logging in to them individually. The gateway APIs enable you to remain logged in to Unified Manager and manage the ONTAP clusters by running ONTAP REST API operations.



All users can run a query by using the GET operation. Application Administrators can run all ONTAP REST operations.

The gateway acts as a proxy to tunnel the API requests by maintaining the header and body requests in the same format as in the ONTAP APIs. You can use your Unified Manager credentials and execute the specific operations to access and manage the ONTAP clusters without passing individual cluster credentials. It continues to manage the cluster authentication and cluster management, but redirects the API requests to run directly on the specific cluster. The response returned by the APIs is the same as the response returned by the respective ONTAP REST APIs executed directly from ONTAP.

HTTP verb	Path (URL)	Description
GET	/gateways	<p>This GET method retrieves the list of all the clusters managed by Unified Manager that support ONTAP REST calls. You can verify the cluster details and choose to run other methods based on the cluster UUID or universal unique identifier (UUID).</p> <p> The gateway APIs retrieve only those clusters supported by ONTAP 9.5 or later, and added to Unified Manager over HTTPS.</p>

HTTP verb	Path (URL)	Description
GET	/gateways/{uuid}/{path}	This is a single point proxy API, supporting POST, DELETE, PATCH operations and GET for all the ONTAP REST APIs. No restrictions apply on any of the API as long as it is supported by ONTAP. The tunnelling or proxy functionality cannot be disabled.
POST		
DELETE		
PATCH		
OPTIONS (not available on Swagger)		
HEAD (not available on Swagger)		The value for {uuid} must be replaced with the cluster UUID on which the REST operation is to be performed. Also, ensure that the UUID is of the cluster supported by ONTAP 9.5 or later, and added to Unified Manager over HTTPS. {path} must be replaced by the ONTAP REST URL. You must remove /api/ from the URL.

## Understanding the API Gateway tunneling

The gateway APIs enable you to manage ONTAP objects through Unified Manager. Unified Manager manages the clusters and authentication details and redirects the requests to the ONTAP REST endpoint. The gateway API transforms the URL and Hypermedia as the Engine of Application State (HATEOAS) links in the header and response body with the API gateway base URL. The gateway API acts as the proxy base URL to which you append the ONTAP REST URL and execute the required ONTAP REST endpoint.

In this example, the gateway API (proxy base URL) is: /gateways/{uuid}/

The ONTAP API taken is: /storage/volumes. You need to add the ONTAP API REST URL as the value for the path parameter.



While adding the path, ensure that you have removed the "/" symbol at the beginning of the URL. For the API /storage/volumes, add storage/volumes.

The appended URL is: /gateways/{uuid}/storage/volumes

On running the GET operation, the generated URL is the following:

GET https://<hostname>/api/gateways/<cluster\_UUID>/storage/volumes

The `/api` tag of the ONTAP REST URL is removed in the appended URL and that for the gateway API is retained.

### Sample cURL command

```
curl -X GET "https://<hostname>/api/gateways/1cd8a442-86d1-11e0-ae1c-9876567890123/storage/volumes" -H "accept: application/hal+json" -H "Authorization: Basic <Base64EncodedCredentials>"
```

The API returns the list of storage volumes in that cluster. The response format is the same as you receive when you run the same API from ONTAP. The status codes returned are the ONTAP REST status codes.

### Setting API scope

All APIs have a context set within the scope of the cluster. APIs that operate on the basis of storage VMs also have the cluster as the scope, that is, the API operations are performed on a particular storage VM within a managed cluster. When you run the `/gateways/{uuid}/{path}` API, ensure that you enter the cluster UUID (Unified Manager datasource UUID) for the cluster on which you run the operation. For setting the context to a particular storage VM within that cluster, enter the storage VM key as the `X-Dot-SVM-UUID` parameter or the storage VM name as the `X-Dot-SVM-Name` parameter. The parameter is added as the filter in the string header and the operation is run within the scope of that storage VM inside that cluster.

### Sample cURL command

```
curl -X GET "https://<hostname>/api/gateways/e4f33f90-f75f-11e8-9ed9-00a098e3215f/storage/volume" -H "accept: application/hal+json" -H "X-Dot-SVM-UUID: d9c33ec0-5b61-11e9-8760-00a098e3215f" -H "Authorization: Basic <Base64EncodedCredentials>"
```

For more information about using ONTAP REST APIs, see [ONTAP REST API Automation](#).

## Performing administrative tasks

You can use the APIs under the `administration` category to modify backup settings, verify the backup file information and cluster certificates, and also manage ONTAP clusters as Active IQ Unified Manager datasources.



You must have the Application Administrator role for running these operations. You can also use the Unified Manager web UI for configuring these settings.

HTTP verb	Path	Description
GET  PATCH	/admin/backup-settings  /admin/backup-settings	<p>You can use the <code>GET</code> method to view the settings of the backup schedule configured in Unified Manager by default. You can verify the following:</p> <ul style="list-style-type: none"> <li>• Whether the schedule is enabled or disabled</li> <li>• Frequency of the backup scheduled (daily or weekly)</li> <li>• Time of the backup</li> <li>• Maximum number of backup files that should be retained in the application</li> </ul> <p>The time of the backup is in server time zone.</p> <p>The database backup settings are available on Unified Manager by default, and you cannot create a backup schedule. However, you can use the <code>PATCH</code> method to modify the default settings.</p>
GET	/admin/backup-file-info	A backup dump file is generated every time the backup schedule is modified for Unified Manager. You can use this method to verify whether the backup file is generated according to the modified backup settings, and whether the information on the file matches the modified settings.
GET	/admin/datasource-certificate	You can use this method to view the datasource (cluster) certificate from the trust store. Validating the certificate is required before adding an ONTAP cluster as a Unified Manager datasource.

HTTP verb	Path	Description
GET	/admin/datasources/clusters	You can use the GET method to retrieve the details of the datasources (ONTAP clusters) managed by Unified Manager.
POST	/admin/datasources/clusters/{key}	
PATCH		You can also add a new cluster to Unified Manager as a datasource. For adding a cluster, you must know its host name, user name, and password.
DELETE		For modifying and deleting a cluster managed as a datasource by Unified Manager, use the ONTAP cluster key.

## Managing users

You can use the APIs in the `security` category to control user access to selected cluster objects in Active IQ Unified Manager. You can add local users or database users. You can also add remote users or groups that belong to an authentication server. Based on the privileges of the roles that you assign to the users, they can manage the storage objects or view the data in Unified Manager.



You must have the Application Administrator role for running these operations. You can also use the Unified Manager web UI for configuring these settings.

The APIs under the `security` category use the `users` parameter, that is the user name, and not the `key` parameter as the unique identifier for the user entity.

HTTP verb	Path	Description
GET	/security/users	You can use these methods to get the details of users or add new user to Unified Manager.
POST	/security/users	You can add specific roles to the users based on their user types. While adding users, you must provide passwords for the local user, maintenance user, and database user.

HTTP verb	Path	Description
GET	/security/users/{name}	
PATCH		
DELETE		

## Viewing performance metrics

Active IQ Unified Manager provides you with a set of APIs under the `/datacenter` category that enable you to view the performance data of your clusters and storage objects in a data center. These APIs retrieve the performance data of the different storage objects, such as clusters, nodes, LUNs, volumes, aggregates, and storage VMs.

The `/metrics` and `/analytics` APIs give different views of the performance metrics, using which, you can drill down to different levels of details for the following storage objects in your data center:

- clusters
- nodes
- storage VMs
- aggregates
- volumes
- LUNs

The following table draws a comparison between the `/metrics` and `/analytics` APIs with respect to the particulars of the performance data retrieved.

Metrics	Analytics
Performance details for a single object. For example, the <code>/datacenter/cluster/clusters/{key}/metrics</code> API requires the cluster key to be entered as the path parameter for retrieving the metrics for that specific cluster.	Performance details for multiple objects of the same type in a data center. For example, the <code>/datacenter/cluster/clusters/analytics</code> API retrieves the collective metrics of all the clusters in a data center.
Performance metrics sample for a storage object based on the time interval parameter for retrieval.	High-level aggregated value of performance for a certain type of storage object for a certain period (above 72 hours).
Basic details of the object are retrieved, such as details of a node or cluster.	No specific details are retrieved.

Metrics	Analytics
<p>Accumulated counters, such as, minimum, maximum, and the average performance values over a period of time are retrieved for a single object, such as, read, write, total, and other counters. For example, the <code>/datacenter/cluster/nodes/{key}/metrics</code> API retrieves the following details (among others) for a node :</p> <pre data-bbox="169 466 768 1833"> {   "iops": {     "local": {       "others": 100.53,       "read": 100.53,       "total": 100.53,       "write": 100.53     },     "others": 100.53,     "read": 100.53,     "total": 100.53,     "write": 100.53   },   "latency": {     "others": 100.53,     "read": 100.53,     "total": 100.53,     "write": 100.53   },   "performance_capacity": {     "available_iops_percent": 0,     "free_percent": 0,     "system_workload_percent": 0,     "used_percent": 0,     "user_workload_percent": 0   },   "throughput": {     "other": 100.53,     "read": 100.53,     "total": 100.53,     "write": 100.53   }, } </pre>	<p>A single aggregated value is displayed for all the objects of the same type. For example, the <code>/datacenter/cluster/nodes/analytics</code> API retrieves the following values (among others) for all the nodes:</p> <pre data-bbox="850 403 1432 882"> {   "iops": 2757.5723,   "latency": 0.8266,   "throughput": 13254.6992,   "utilization_percent": 58.6488,   "period": 72,   "performance_capacity": {     "used_percent": 63.5317,     "available_iops_percent": 60731.3086   }, } </pre>

Metrics	Analytics
<p>The time range and sample data is based on the following schedule: The time range for the data. Examples can be 1h, 12h, 1d, 2d, 3d, 15d, 1w, 1m, 2m, 3m, 6m. You get 1 hour samples if the range is more than 3 days (72 hr) else it is 5 minutes samples. The period for each time range is as follows:</p> <ul style="list-style-type: none"> <li>• 1h: Metrics over the most recent hour sampled over 5 minutes.</li> <li>• 12h: Metrics over the most recent 12 hours sampled over 5 minutes.</li> <li>• 1d: Metrics over the most recent day sampled over 5 minutes.</li> <li>• 2d: Metrics over the most recent 2 days sampled over 5 minutes.</li> <li>• 3d: Metrics over the most recent 3 days sampled over 5 minutes.</li> <li>• 15d: Metrics over the most recent 15 days sampled over 1 hour.</li> <li>• 1w: Metrics over the most recent week sampled over 1 hour.</li> <li>• 1m: Metrics over the most recent month sampled over 1 hour.</li> <li>• 2m: Metrics over the most recent 2 months sampled over 1 hour.</li> <li>• 3m: Metrics over the most recent 3 months sampled over 1 hour.</li> <li>• 6m: Metrics over the most recent 6 months sampled over 1 hour.</li> </ul> <p>Available values : 1h, 12h, 1d, 2d, 3d, 15d, 1w, 1m, 2m, 3m, 6m</p> <p>Default value : 1h</p>	<p>Above 72 hours, that is 3 days. The duration over which this sample is calculated are represented in the ISO-8601 standard format.</p>

The following table describes the /metrics and /analytics APIs in details.



The IOPS and performance metrics returned by these APIs are double values, for example 100.53. Filtering these float values by the pipe (|) and wildcard (\*) characters is not supported.

HTTP Verb	Path	Description
GET	/datacenter/cluster/clusters/{key}/metrics	Retrieves performance data (sample and summary) for a cluster specified by the input parameter of the cluster key. Information, such as the cluster key and UUID, time range, IOPS, throughput, and number of samples is returned.
GET	/datacenter/cluster/clusters/analytics	Retrieves high-level performance metrics for all the clusters in a data center. You can filter your results based on the required criteria. Values, such as aggregated IOPS, throughput, and the period of collection (in hours) are returned.
GET	/datacenter/cluster/nodes/{key}/metrics	Retrieves performance data (sample and summary) for a node specified by the input parameter of the node key. Information, such as the node UUID, time range, summary of the IOPS, throughput, latency, and performance, the number of samples collected, and percentage utilized is returned.
GET	/datacenter/cluster/nodes/analytics	Retrieves high-level performance metrics for all the nodes in a data center. You can filter your results based on the required criteria. Information, such as node and cluster keys, and values, such as aggregated IOPS, throughput, and the period of collection (in hours) are returned.
GET	/datacenter/storage/aggregates/{key}/metrics	Retrieves performance data (sample and summary) for an aggregate specified by the input parameter of the aggregate key. Information, such as the time range, summary of the IOPS, latency, throughput, and performance capacity, the number of samples collected for each counter, and percentage utilized is returned.

HTTP Verb	Path	Description
GET	/datacenter/storage/aggregates/analytics	Retrieves high-level performance metrics for all the aggregates in a data center. You can filter your results based on the required criteria. Information, such as aggregate and cluster keys, and values, such as aggregated IOPS, throughput, and the period of collection (in hours) are returned.
GET	/datacenter/storage/luns/{key}/metrics /datacenter/storage/volumes/{key}/metrics	Retrieves performance data (sample and summary) for a LUN or a file share (volume) specified by the input parameter of the LUN or volume key. Information, such as the summary of the minimum, maximum, and average of the read, write, and total IOPS, latency, and throughput, and the number of samples collected for each counter is returned.
GET	/datacenter/storage/luns/analytics /datacenter/storage/volumes/analytics	Retrieves high-level performance metrics for all the LUNs or volumes in a data center. You can filter your results based on the required criteria. Information, such as storage VM and cluster keys, and values, such as aggregated IOPS, throughput, and the period of collection (in hours) are returned.
GET	/datacenter/svm/svms/{key}/metrics	Retrieves performance data (sample and summary) for a storage VM specified by the input parameter of the storage VM key. Summary of the IOPS based on each supported protocol, such as nvmf, fcp, iscsi, and nfs, throughput, latency, and the number of samples collected are returned.

HTTP Verb	Path	Description
GET	/datacenter/svm/svms/analytics	Retrieves high-level performance metrics for all the storage VMs in a data center. You can filter your results based on the required criteria. Information, such as storage VM UUID, aggregated IOPS, latency, throughput, and the period of collection (in hours) are returned.
GET	/datacenter/cluster/licensing/licenses /datacenter/cluster/licensing/licenses/{key}	Returns the details of the licenses installed on the clusters in your data center. You can filter your results based on the required criteria. Information, such as the license key, cluster key, expiry date, and license scope is returned. You can enter a license key to retrieve the details of a specific license.

## Viewing Jobs

You can use the `jobs` API under the `management-server` category to view the execution details of asynchronous operations.

In Active IQ Unified Manager, operations, such as adding and modifying resources are performed by synchronous and asynchronous API invocations. Invocations that are scheduled for asynchronous execution can be tracked by a Job object created for that invocation. Each Job object has a unique key for identification. Each Job object returns the Job object URI for you to access and track the progress of the job. You can use this API for retrieving the details of each execution.

By using this API, you can query all the Job objects for your data center, including historical data. Querying all the jobs, by default, returns the details of the last 20 jobs triggered through the web UI and API interface. Use the inbuilt filters to view specific Jobs. You can also use the Job key to query the details of a specific job and run the next set of operations on the resources.

Category	HTTP verb	Path	Description
management-server	GET	/management-server/jobs	Returns the job details of all the jobs. Without any sort order, the last submitted Job object is returned on top.

Category	HTTP verb	Path	Description
management-server	GET	/management-server/jobs/ { key}  Enter the job key of the Job object to view the specific details of that job.	Returns the details of the specific Job object.

## Viewing events and system details

The events and system APIs under the management-server category enable you to retrieve the events that are generated for the monitored clusters in your data center and view the instance details in your Active IQ Unified Manager environment respectively.

### Viewing events

By using the /management-server/events API, you can query the events in your data center, including historical data. Use the inbuilt filters, such as name, impact level, impact area, severity, state, resource name and resource type, to view specific events. The resource type and area parameters return information about the storage object on which the event has occurred, and the impact area returns the information about the issue for which the event is raised, such as availability, capacity, configuration, security, protection and performance.

You can also use the event key to query the details of a specific event and run the next set of operations on the resources.

Category	HTTP verb	Path	Description
management-server	GET	/management-server/events	The response body consists of the event details of the queried events in that data center.
management-server	GET	/management-server/events/ { key}	Run this API if you want to query a particular event. Enter the event key of the job to view the details. The response body consists of the details of that event.

### Viewing system details

By using the /management-server/system API, you can query the instance-specific details of your Unified Manager environment. The API returns information about the product and services, such as the version of Unified Manager installed on your system, UUID, vendor name, host OS, and the name, description, and status of the services running on the Unified Manager instance.

Category	HTTP verb	Path	Description
management-server	GET	/management-server/system	No input parameter is required for running this API. The system details of the current Unified Manager instance are returned by default.

## Managing workloads

The APIs described here cover various functions of storage administration, such as viewing storage workloads, creating LUNs and file shares, managing Performance Service Levels and Storage Efficiency Policies, and assigning the policies on storage workloads.

### Viewing storage workloads

The APIs listed here enable you to view a consolidated list of storage workloads for all of the ONTAP clusters in your data center. The APIs also provide a summary view of the number of the storage workloads provisioned in your Active IQ Unified Manager environment, and their capacity and performance (IOPS) statistics.

#### View storage workloads

You can use the following method to view all the storage workloads in all the clusters in your data center. For information about filtering the response based on specific columns, see the API reference documentation available in your Unified Manager instance.

Category	HTTP verb	Path
storage-provider	GET	/storage-provider/workloads

#### View storage workloads summary

You can use the following method to assess the used capacity, available capacity, used IOPS, available IOPS, and number of storage workloads managed by each Performance Service Level. The storage workloads displayed can be for any LUN, NFS file share, or CIFS share. The API provides a storage workloads overview, an overview of the storage workloads provisioned by the Unified Manager, a data center overview, an overview of the total, used, and available space and IOPS in the data center, in terms of the assigned Performance Service Levels. The information received in response to this API is used to populate the dashboard in the Unified Manager UI.

Category	HTTP verb	Path
storage-provider	GET	/storage-provider/workloads-summary

## Managing access endpoints

You need to create access endpoints or logical interfaces (LIFs), which are required for provisioning Storage Virtual Machines (SVMs), LUNs, and file shares. You can view, create, modify, and delete the access endpoints for the SVMs, LUNs, or file shares in your Active IQ Unified Manager environment.

### View access endpoints

You can view a list of the access endpoints in your Unified Manager environment by using the following method. To query a list of access endpoints of a particular SVM, LUN, or file share, you need to enter the unique identifier for the SVM, LUN, or file share. You can also enter the unique access endpoint key to retrieve the details of the particular access endpoint.

Category	HTTP verb	Path
storage-provider	GET	/storage-provider/access-endpoints /storage-provider/access-endpoints/{key}

### Add access endpoints

You can create custom access endpoints and assign required properties to it. You must enter the details of the access endpoint that you want to create as the input parameters. You can use this API, or the System Manager or ONTAP CLI to create an access endpoint on each node. Both IPv4 and IPv6 addresses are supported for access endpoints creation.

 You must configure your SVM with a minimum number of access endpoints per node for successful provisioning of LUNs and file shares. You should configure your SVM with at least two access endpoints per node, one supporting CIFS and/or NFS protocol, another supporting iSCSI or FCP protocol.

Category	HTTP verb	Path
storage-provider	POST	/storage-provider/access-endpoints

### Delete access endpoints

You can delete a specific access endpoint by using the following method. You need to provide the access endpoint key as an input parameter to delete a particular access endpoint.

Category	HTTP verb	Path
storage-provider	DELETE	/storage-provider/access-endpoints/{key}

## Modify access endpoints

You can modify an access endpoint and update its properties by using the following method. You need to provide the access endpoint key to modify a particular access endpoint. You also need to enter the property that you want to update, along with its value.

Category	HTTP verb	Path
storage-provider	PATCH	/storage-provider/access-endpoints/{key}

## Managing Active Directory mapping

You can use the APIs listed here to manage Active Directory mappings on the SVM that are required for provisioning CIFS shares on the SVMs. Active Directory mappings need to be configured for mapping the SVMs with ONTAP.

### View Active Directory mappings

You can view the configuration details of the Active Directory mappings for an SVM by using the following method. For viewing the Active Directory mappings on an SVM, you need to enter the SVM key. For querying the details of a particular mapping, you must enter the mapping key.

Category	HTTP verb	Path
storage-provider	GET	/storage-provider/active-directories-mappings
		/storage-provider/active-directories-mappings/{key}

### Add Active Directory mapping

You can create Active Directory mappings on an SVM by using the following method. You must enter the mapping details as the input parameters.

Category	HTTP verb	Path
storage-provider	POST	/storage-provider/active-directories-mappings

## Managing file shares

You can use the `/storage-provider/file-shares` API to view, add, modify, and delete the CIFS and NFS file share volumes in your data center environment.

Before provisioning the file share volumes, ensure that the SVM has been created and provisioned with the supported protocols. If you are assigning Performance Service Levels (PSLs) or Storage Efficiency Policies (SEPs), while provisioning, the PSLs or SEPs should be created before creating the file shares.

## View file shares

You can use the following method to view the file share volumes available in your Unified Manager environment. When you have added an ONTAP cluster as a datasource on Active IQ Unified Manager, the storage workloads for those clusters are automatically added to your Unified Manager instance. This API retrieves the file shares automatically and manually added to your Unified Manager instance. You can view the details of a specific file share by running this API with the file share key.

Category	HTTP verb	Path
storage-provider	GET	/storage-provider/file-shares /storage-provider/file-shares/{key}

## Add file shares

You can use the following method to add CIFS and NFS file shares in your SVM. You must enter the details of the file share that you want to create, as the input parameters. You cannot use this API for adding FlexGroup volumes.

Category	HTTP verb	Path
storage-provider	POST	/storage-provider/file-shares



Depending on whether the access control list (ACL) parameters or the export policy parameters are provided, CIFS shares or NFS file shares are created. If you do not provide the values for the ACL parameters, CIFS shares are not created, and NFS shares are created by default, providing access to all.

**Creating data-protection volumes:** When you add file shares to your SVM, the type of the volume that is mounted, by default, is `rw` (read-write). For creating data-protection (DP) volumes, specify `dp` as the value for the `type` parameter.

## Delete file shares

You can use the following method to delete a specific file share. You need to enter the file share key as an input parameter to delete a particular file share.

Category	HTTP verb	Path
storage-provider	DELETE	/storage-provider/file-shares/{key}

## Modify file shares

You can use the following method to modify a file share and update its properties.

You need to provide the file share key to modify a particular file share. Additionally, you need to enter the

property that you want to update, along with its value.



Note that you can update only one property at a single invocation of this API. For multiple updates, you need to run this API as many times.

Category	HTTP verb	Path
storage-provider	PATCH	/storage-provider/file-shares/{key}

## Managing LUNs

You can use the `/storage-provider/luns` API to view, add, modify, and delete the LUNs in your data center environment.

Before provisioning the LUNs, ensure that the SVM has been created and provisioned with the supported protocols. If you are assigning Performance Service Levels (PSLs) or Storage Efficiency Policies (SEPs), while provisioning, the PSLs or SEPs should be created before creating the LUN.

### View LUNs

You can use the following method to view the LUNs in your Unified Manager environment. When you have added an ONTAP cluster as a datasource on Active IQ Unified Manager, the storage workloads for those clusters are automatically added to your Unified Manager instance. This API retrieves all the LUNs automatically and manually added to your Unified Manager instance. You can view the details of a specific LUN by running this API with the LUN key.

Category	HTTP verb	Path
storage-provider	GET	/storage-provider/luns /storage-provider/luns/{key}

### Add LUNs

You can use the following method to add LUNs to your SVMs.

Category	HTTP verb	Path
storage-provider	POST	/storage-provider/luns



In your cURL request, if you provide a value for the optional parameter `volume_name_tag` in the input, then that value is used while naming the volume during the LUN creation. This tag enables searching the volume easily. If you provide the `volume` key in the request, the tagging is skipped.

## Delete LUNs

You can use the following method to delete a specific LUN. You need to provide the LUN key to delete a particular LUN.



If you have created a volume in ONTAP and then provisioned LUNs through Unified Manager on that volume, when you delete all of the LUNs by using this API, the volume also gets deleted from the ONTAP cluster.

Category	HTTP verb	Path
storage-provider	DELETE	/storage-provider/luns/{key}

## Modify LUNs

You can use the following method to modify a LUN and update its properties. You need to provide the LUN key to modify a particular LUN. You also need to enter the LUN property that you want to update, along with its value. For updating LUN arrays by using this API, you should review the recommendations in “Recommendations for using the APIs”.



You can update only one property at a single invocation of this API. For multiple updates, you need to run this API as many times.

Category	HTTP verb	Path
storage-provider	PATCH	/storage-provider/luns/{key}

## Managing Performance Service Levels

You can view, create, modify, and delete Performance Service Levels by using the storage provider APIs for on your Active IQ Unified Manager.

### View Performance Service Levels

You can use the following method to view the Performance Service Levels for assigning them to storage workloads. The API lists all of the system-defined and user-created Performance Service Levels, and retrieves the attributes of all of the Performance Service Levels. If you want to query a specific Performance Service Level, you need to enter the unique ID of the Performance Service Level to retrieve its details.

Category	HTTP verb	Path
storage-provider	GET	/storage-provider/performance-service-levels  /storage-provider/performance-service-levels/{key}

## Add Performance Service Levels

You can use the following method to create custom Performance Service Levels and assign them to your storage workloads if the system-defined Performance Service Levels do not meet the required service level objectives (SLOs) for the storage workloads. Enter the details for the Performance Service Level that you want to create. For the IOPS properties, ensure that you enter valid range of values.

Category	HTTP verb	Path
storage-provider	POST	/storage-provider/performance-service-levels

## Delete Performance Service Levels

You can use the following method to delete a specific Performance Service Level. You cannot delete a Performance Service Level if it is assigned to a workload, or if it is the only available Performance Service Level. You need to provide the unique ID of the Performance Service Level as an input parameter to delete a particular Performance Service Level.

Category	HTTP verb	Path
storage-provider	DELETE	/storage-provider/performance-service-levels/{key}

## Modify Performance Service Levels

You can use the following method to modify a Performance Service Level and update its properties. You cannot modify a Performance Service Level that is system-defined or is assigned to a workload. You need to provide the unique ID of the to modify a particular Performance Service Level. You must also enter the IOPS property that you want to update, along with a valid value.

Category	HTTP verb	Path
storage-provider	PATCH	/storage-provider/performance-service-levels/{key}

## Viewing aggregate capabilities based on Performance Service Levels

You can use the following method to query the aggregate capabilities based on Performance Service Levels. This API returns the list of aggregates available in your data center and indicates the capabilities in terms of the Performance Service Levels that can be supported in those aggregates. While provisioning workloads on a volume, you can view the capability of an aggregate to support a particular Performance Service Level, and provision workloads based on that capability. Your ability to specify the aggregate is available only when you are provisioning a workload by using APIs. This functionality is not available on the Unified Manager web UI.

Category	HTTP verb	Path
storage-provider	GET	/storage-provider/aggregate-capabilities  /storage-provider/aggregate-capabilities/{key}

## Managing Storage Efficiency Policies

You can view, create, modify, and delete Storage Efficiency Policies by using the storage provider APIs.

Note the following points:

- It is not mandatory to assign a Storage Efficiency Policy while creating a workload on Unified Manager.
- You cannot unassign a Storage Efficiency Policy from a workload after a policy is assigned to it.
- If a workload has some storage settings specified on ONTAP volumes, such as deduplication and compression, those settings can be overwritten by the settings specified in the Storage Efficiency Policy that you apply when you add the storage workloads on Unified Manager.

### View Storage Efficiency Policies

You can use the following method to view the Storage Efficiency Policies before assigning them to storage workloads. This API lists all of the system-defined and user-created Storage Efficiency Policies, and retrieves the attributes of all of the Storage Efficiency Policies. If you want to query a specific Storage Efficiency Policy, you need to enter the unique ID of the policy to retrieve its details.

Category	HTTP verb	Path
storage-provider	GET	/storage-provider/storage-efficiency-policies  /storage-provider/storage-efficiency-policies/{key}

## Add Storage Efficiency Policies

You can use the following method to create custom Storage Efficiency Policies, and assign them to your storage workloads if the system-defined policies do not meet the provisioning requirements for your storage workloads. Enter the details of the Storage Efficiency Policy that you want to create, as input parameters.

Category	HTTP verb	Path
storage-provider	POST	/storage-provider/storage-efficiency-policies

## Delete Storage Efficiency Policies

You can use the following method to delete a specific Storage Efficiency Policy. You cannot delete a Storage Efficiency Policy if it is assigned to a workload, or if it is the only available Storage Efficiency Policy. You need to provide the unique ID of the Storage Efficiency Policy as an input parameter to delete a particular Storage Efficiency Policy.

Category	HTTP verb	Path
storage-provider	DELETE	/storage-provider/storage-efficiency-policies/{key}

## Modify Storage Efficiency Policies

You can use the following method to modify a Storage Efficiency Policy and update its properties. You cannot modify a Storage Efficiency Policy that is system-defined or is assigned to a workload. You need to provide the unique ID of the Storage Efficiency Policy to modify a particular Storage Efficiency Policy. Additionally, you need to provide the property that you want to update, along with its value.

Category	HTTP verb	Path
storage-provider	PATCH	/storage-provider/storage-efficiency-policies/{key}

## Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

**LIMITED RIGHTS LEGEND:** Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.