



Manage storage using REST APIs

Active IQ Unified Manager 9.10

NetApp

November 12, 2024

Table of Contents

- Manage storage using REST APIs 1
 - Getting started with Active IQ Unified Manager REST APIs 1
 - REST API access and authentication in Active IQ Unified Manager 5
 - Unified Manager REST APIs 15
 - Common workflows for storage management 47

Manage storage using REST APIs

Getting started with Active IQ Unified Manager REST APIs

Active IQ Unified Manager provides a set of APIs to manage your storage resources on the supported storage systems through a RESTful web service interface for any third-party integration.

In these topics, you will find information about Unified Manager APIs, sample workflows to resolve specific issues, and some sample codes. Using this information, you can create RESTful clients of NetApp Manageability software solutions for managing NetApp systems. The APIs are based on the Representational State Transfer (REST) architectural style. All the four REST operations Create, Read, Update, and Delete (also known as CRUD) are supported.

Audience for this content

The topics here are intended for developers creating applications that interface with the Active IQ Unified Manager software through REST APIs.

Storage administrators and architects, can refer to this information to gain a basic understanding of how the Unified Manager REST APIs can be used to build client applications to manage and monitor NetApp storage systems.

You should use this information if you want to use the storage provider, ONTAP cluster, and management administration APIs for managing your storage.



You must have one of the following roles: Operator, Storage Administrator, or Application Administrator. You must know the IP address or fully qualified domain name of the Unified Manager server on which you want to execute the REST APIs.

Active IQ Unified Manager API access and categories

The Active IQ Unified Manager APIs enable you to manage and provision storage objects in your environment. You can also access the Unified Manager web UI to perform some of these functions.

Constructing a URL to directly access REST APIs

You can access the REST APIs directly through a programming language, such as Python, C#, C++, JavaScript, and so forth. Enter the host name or IP address and the URL to access the REST APIs in the format

`https://<hostname>/api`



The default port is 443. You can configure the port as required by your environment.

Accessing the online API documentation page

You can access the *API Documentation* reference content page that is packaged along with the product to display the API documentation, as well as to manually issue an API call (on the interface, for example,

Swagger). You can access this documentation on clicking the **Menu Bar > Help button > API Documentation**

Alternatively, enter the host name or IP address and the URL to access the REST API page in the format

`https://<hostname>/docs/api/`

Categories

The API calls are organized into functionally based on the areas or categories. To locate a specific API, click the applicable API category.

The REST APIs provided with Unified Manager help you to perform administrative, monitoring, and provisioning functions. The APIs are grouped under the following categories.

- **datacenter**

This category contains the APIs that help you in datacenter storage management and analytics using tools, such as Work Flow Automation and Ansible. The REST APIs under this category provide information about the clusters, nodes, aggregates, volumes, LUNs, file shares, namespaces, and other elements in your data center.

- **management-server**

The APIs under the **management-server** category contain the `jobs`, `system`, and `events` APIs. Jobs are operations that are scheduled for asynchronous execution related to managing of storage objects or workloads on Unified Manager. The `events` API returns events in your data center, and the `system` API returns the Unified Manager instance details.

- **storage-provider**

This category contains all of the provisioning APIs required for managing and provisioning file shares, LUNs, Performance Service Levels, and Storage Efficiency Policies. The APIs also enable you to configure access endpoints, Active Directories, as well as assign Performance Service Levels and Storage Efficiency Policies on storage workloads.

- **administration**

This category contains the APIs used for running administrative tasks, such as maintaining backup settings, viewing trust store certificates for the Unified Manager datasources, and managing ONTAP clusters as datasources for Unified Manager.

- **gateway**

Unified Manager enables you to invoke ONTAP REST APIs through the APIs under the gateway category and manage the storage objects in your data center.

- **security**

This category contains APIs for managing Unified Manager users.

REST services offered in Active IQ Unified Manager

You should be aware of the REST services and operations offered, before you start using the Active IQ Unified Manager APIs.

The provisioning and administrative APIs that are used for configuring the API server support the read (GET) or write (POST, PATCH, DELETE) operations. The following are some examples of the GET, PATCH, POST, and DELETE operations that are supported by the APIs:

- Example for GET: `GET /datacenter/cluster/clusters` retrieves cluster details in your data center. The maximum number of records that is returned by the GET operation is 1000.



The APIs enable you to filter, sort, and order the records by supported attributes.

- Example for POST: `POST /datacenter/svm/svms` creates a custom Storage Virtual Machine (SVM).
- Example for PATCH: `PATCH /datacenter/svm/svms/{key}` modifies the properties of an SVM, using its unique key.
- Example for DELETE: `DELETE /storage-provider/access-endpoints/{key}` deletes an access endpoint from a LUN, SVM, or file share by using its unique key.

The REST operations that can be performed by using the APIs depend on the role of the Operator, Storage Administrator, or Application Administrator user.

User role	Supported REST method
Operator	Read-only access to data. Users with this role can run all GET requests.
Storage Administrator	Read access to all data. Users with this role can run all GET requests. Additionally, they have write access (to run PATCH, POST, and DELETE requests) to perform specific activities, such as managing, storage service objects, and storage management options.
Application Administrator	Read and write access to all data. Users with this role can run GET, PATCH, POST, and DELETE requests for all functions.

For more information about all the REST operations, see the *Online API documentation*.

API version in Active IQ Unified Manager

The REST API URIs in Active IQ Unified Manager specifies a version number. For example, `/v2/datacenter/svm/svms`. The version number `v2` in `/v2/datacenter/svm/svms` indicates the API version used in a specific release. The version number minimizes the impact of API changes on the client software by sending back a response that the client can process.

The numerical part of this version number is incremental with respect to releases. URIs with a version number provide a consistent interface that maintains backward compatibility in future releases. You also find the same APIs without a version, for example `/datacenter/svm/svms`, that indicate the base APIs without a version. The base APIs are always the latest version of the APIs.



On the top right corner of your Swagger interface, you can select the version of the API to use. The highest version is selected by default. It is recommended that you use the highest version of a particular API (with respect to the incremental integer) available in your Unified Manager instance.

For all requests, you must explicitly request the API version that you want to use. When the version number is specified, the service does not return response elements that your application is not designed to handle. In REST requests, you should include the version parameter. The earlier versions of the APIs are eventually deprecated after a few releases. In this release, the v1 version of the APIs is deprecated.

Storage resources in ONTAP

The storage resources in ONTAP can be broadly classified into *physical storage resources* and *logical storage resources*. To effectively manage your ONTAP systems using the APIs provided in Active IQ Unified Manager, you must understand the storage resource model and the relationship between various storage resources.

- **Physical storage resources**


Refers to the physical storage objects provided by ONTAP. Physical storage resources include disks, clusters, storage controllers, nodes, and aggregates.

- **Logical storage resources**

Refers to the storage resources provided by ONTAP that are not tied to a physical resource. These resources are associated with a storage virtual machine (SVM, formerly known as Vserver), and they exist independently of any specific physical storage resource such as a disk, array LUN, or aggregate.

Logical storage resources include volumes of all types and qtrees, as well as the capabilities and configurations you can use with these resources, such as Snapshot copies, deduplication, compression, and quotas.

The following illustration shows the storage resources in a 2-node cluster:



REST API access and authentication in Active IQ Unified Manager

The Active IQ Unified Manager REST API is accessible by using any REST client or programming platform that can issue HTTP requests with a basic HTTP authentication mechanism.

A sample request and response:

- **Request**

```
GET
https://<IP
address/hostname>:<port_number>/api/v2/datacenter/cluster/clusters
```

- **Response**

```
{
  "records": [
    {
      "key": "4c6bf721-2e3f-11e9-a3e2-
```

```

00a0985badbb:type=cluster,uuid=4c6bf721-2e3f-11e9-a3e2-00a0985badbb",
  "name": "fas8040-206-21",
  "uuid": "4c6bf721-2e3f-11e9-a3e2-00a0985badbb",
  "contact": null,
  "location": null,
  "version": {
    "full": "NetApp Release Dayblazer__9.5.0: Thu Jan 17 10:28:33
UTC 2019",
    "generation": 9,
    "major": 5,
    "minor": 0
  },
  "isSanOptimized": false,
  "management_ip": "10.226.207.25",
  "nodes": [
    {
      "key": "4c6bf721-2e3f-11e9-a3e2-
00a0985badbb:type=cluster_node,uuid=12cf06cc-2e3a-11e9-b9b4-
00a0985badbb",
      "uuid": "12cf06cc-2e3a-11e9-b9b4-00a0985badbb",
      "name": "fas8040-206-21-01",
      "_links": {
        "self": {
          "href": "/api/datacenter/cluster/nodes/4c6bf721-2e3f-11e9-
a3e2-00a0985badbb:type=cluster_node,uuid=12cf06cc-2e3a-11e9-b9b4-
00a0985badbb"
        }
      },
      "location": null,
      "version": {
        "full": "NetApp Release Dayblazer__9.5.0: Thu Jan 17
10:28:33 UTC 2019",
        "generation": 9,
        "major": 5,
        "minor": 0
      },
      "model": "FAS8040",
      "uptime": 13924095,
      "serial_number": "701424000157"
    },
    {
      "key": "4c6bf721-2e3f-11e9-a3e2-
00a0985badbb:type=cluster_node,uuid=1ed606ed-2e3a-11e9-a270-
00a0985bb9b7",
      "uuid": "1ed606ed-2e3a-11e9-a270-00a0985bb9b7",
      "name": "fas8040-206-21-02",

```



```

    "_links": {
      "self": {
        "href": "/api/datacenter/cluster/nodes/4c6bf721-2e3f-11e9-
a3e2-00a0985badbb:type=cluster_node,uuid=1ed606ed-2e3a-11e9-a270-
00a0985bb9b7"
      }
    },
    "location": null,
    "version": {
      "full": "NetApp Release Dayblazer__9.5.0: Thu Jan 17
10:28:33 UTC 2019",
      "generation": 9,
      "major": 5,
      "minor": 0
    },
    "model": "FAS8040",
    "uptime": 14012386,
    "serial_number": "701424000564"
  }
],
"_links": {
  "self": {
    "href": "/api/datacenter/cluster/clusters/4c6bf721-2e3f-11e9-
a3e2-00a0985badbb:type=cluster,uuid=4c6bf721-2e3f-11e9-a3e2-
00a0985badbb"
  }
}
},

```

- *IP address/hostname* is the IP address or the fully qualified domain name (FQDN) of the API server.
- Port 443

443 is the default HTTPS port. You can customize the HTTPS port, if required.

To issue HTTP requests from a web browser, you have to use REST API browser plugins. You can also access the REST API by using scripting platforms such as cURL and Perl.

Authentication

Unified Manager supports the basic HTTP authentication scheme for APIs. For secure information flow (request and response), the REST APIs are accessible only over HTTPS. The API server provides a self-signed SSL certificate to all clients for server verification. This certificate can be replaced by a custom certificate (or CA certificate).

You must configure user access to the API server for invoking the REST APIs. The users can be local users (user profiles stored in the local database) or LDAP users (if you have configured the API server to authenticate over LDAP). You can manage user access by logging in to the Unified Manager Administration

Console user interface.

HTTP status codes used in Active IQ Unified Manager

While running the APIs or troubleshooting issues, you should be aware of the various HTTP status codes and error codes that are used by Active IQ Unified Manager APIs.

The following table lists the error codes related to authentication:

HTTP status code	Status code title	Description
200	OK	Returned on successful execution of synchronous API calls.
201	Created	Creation of new resources by synchronous calls, such as configuration of Active Directory.
202	Accepted	Returned on successful execution of asynchronous calls for provisioning functions, such as creating LUNs and files shares.
400	Invalid request	Indicates input validation failure. User has to correct the inputs, for example, valid keys in a request body.
401	Unauthorized request	You are not authorized to view the resource/Unauthorized.
403	Forbidden request	Accessing the resource you were trying to reach is forbidden.
404	Resource not found	The resource you were trying to reach is not found.
405	Method Not Allowed	Method not allowed.
429	Too Many Requests	Returned when the user sends too many requests within a specific time.

HTTP status code	Status code title	Description
500	Internal server error	Internal server error. Failed to get the response from server. This internal server error may or may not be permanent. For example, if you run a <code>GET</code> or <code>GET ALL</code> operation and receive this error, it is recommended that you repeat this operation for a minimum of five retries. If it is a permanent error, then the status code returned continues to be 500. If the operation succeeds, the status code returned is 200.

Recommendations for using the APIs for Active IQ Unified Manager

When using the APIs in Active IQ Unified Manager, you should follow certain recommended practices.

- All response content type must be in the following format for a valid execution:

```
application/json
```

- The API version number is not related to the product version number. You should use the latest version of the API available for your Unified Manager instance. For more information about Unified Manager API versions, see the “REST API versioning in Active IQ Unified Manager” section.
- While updating array values using a Unified Manager API, you must update the entire string of values. You cannot append values to an array. You can only replace an existing array.
- You can use filter operators, such as pipe (`()`) and wild card (`*`) for all query parameters, except for double values, for example, IOPS and performance in the metrics APIs.
- Avoid querying objects by using a combination of the filter operators wild card (`*`) and pipe (`()`). It might retrieve an incorrect number of objects.
- When using values for filter, ensure that the value does not contain any `?` character. This is to mitigate risks of SQL injection.
- Note that the `GET (all)` request for any API returns a maximum of 1000 records. Even if you run the query by setting the `max_records` parameter to a value higher than 1000, only 1000 records are returned.
- For performing administrative functions, it is recommended that you use the Unified Manager UI.

Logs for troubleshooting

System logs enable you to analyze the causes of failure and troubleshooting issues that may arise while running the APIs.

Retrieve the logs from the following location for troubleshooting issues related to the API calls.

Log location	Use
<code>/var/log/ocie/access_log.log</code>	<p>Contains all API call details, such as the user name of the user invoking the API, start time, execution time, status, and URL.</p> <p>You can use this log file to check the frequently-used APIs, or troubleshoot any GUI workflow. You can also use it to scale analysis, based on the execution time.</p>
<code>/var/log/ocum/ocumserver.log</code>	<p>Contains all API execution logs.</p> <p>You can use this log file to troubleshoot and debug the API calls.</p>
<code>/var/log/ocie/server.log</code>	<p>Contains all Wildfly server deployments and start/stop service related logs.</p> <p>You can use this log file to find the root cause of any issues occurring during the start, stop, or deployment of the Wildfly server.</p>
<code>/var/log/ocie/au.log</code>	<p>Contains acquisition unit related logs.</p> <p>You can use this log file when you have created, modified, or deleted any objects in ONTAP but they do not get reflected for the Active IQ Unified Manager REST APIs.</p>

Job objects asynchronous processes

Active IQ Unified Manager provides the `jobs` API that retrieves information about the Jobs performed while running other APIs. you must know how asynchronous processing works using the Job object.

Some of the API calls, particularly those that are used for adding or modifying resources, can take longer to complete than other calls. Unified Manager processes these long-running requests asynchronously.

Asynchronous requests described using Job object

After making an API call that runs asynchronously, the HTTP response code 202 indicates the request has been successfully validated and accepted, but not yet completed. The request is processed as a background task which continues to run after the initial HTTP response to the client. The response includes the Job object anchoring the request, including its unique identifier.

Querying the Job object associated with an API request

The Job object returned in the HTTP response contains several properties. You can query the state property to determine if the request completed successfully. A Job object can be in one of the following states:

- NORMAL

- WARNING
- PARTIAL_FAILURES
- ERROR

There are two techniques you can use when polling a Job object to detect a terminal state for the task, either success or failure:

- Standard polling request: The current Job state is returned immediately.
- Long polling request: When the job state moves to NORMAL, ERROR, or PARTIAL_FAILURES.

Steps in an asynchronous request

You can use the following high-level procedure to complete an asynchronous API call:

1. Issue the asynchronous API call.
2. Receive an HTTP response 202 indicating successful acceptance of the request.
3. Extract the identifier for the Job object from the response body.
4. Within a loop, wait for the Job object to reach the terminal state NORMAL, ERROR, or PARTIAL_FAILURES.
5. Verify the terminal state of the Job and retrieve the Job result.

Hello API server

The *Hello API server* is a sample program that demonstrates how to invoke a REST API in Active IQ Unified Manager using a simple REST client. The sample program provides you basic details about the API server in the JSON format (the server supports only `application/json` format).

The URI used is: <https://<hostname>/api/datacenter/svm/svms>. This sample code takes the following input parameters:

- The API server IP address or FQDN
- Optional: Port number (default: 443)
- User name
- Password
- Response format (`application/json`)

To invoke REST APIs, you can also use other scripts such as Jersey and RESTEasy to write a Java REST client for Active IQ Unified Manager. You should be aware of the following considerations about the sample code:

- Uses an HTTPS connection to Active IQ Unified Manager to invoke the specified REST URI
- Ignores the certificate provided by Active IQ Unified Manager
- Skips the host name verification during the handshake
- Uses `javax.net.ssl.HttpURLConnection` for a URI connection

- Uses a third-party library (`org.apache.commons.codec.binary.Base64`) for constructing the Base64 encoded string used in the HTTP basic authentication

To compile and execute the sample code, you must use Java compiler 1.8 or later.

```
import java.io.BufferedReader;
import java.io.InputStreamReader;
import java.net.URL;
import java.security.SecureRandom;
import java.security.cert.X509Certificate;
import javax.net.ssl.HostnameVerifier;
import javax.net.ssl.HttpURLConnection;
import javax.net.ssl.SSLContext;
import javax.net.ssl.SSLSession;
import javax.net.ssl.TrustManager;
import javax.net.ssl.X509TrustManager;
import org.apache.commons.codec.binary.Base64;

public class HelloApiServer {

    private static String server;
    private static String user;
    private static String password;
    private static String response_format = "json";
    private static String server_url;
    private static String port = null;

    /*
     * * The main method which takes user inputs and performs the *
    necessary steps
     * to invoke the REST URI and show the response
    */ public static void main(String[] args) {
        if (args.length < 2 || args.length > 3) {
            printUsage();
            System.exit(1);
        }
        setUserArguments(args);
        String serverBaseUrl = "https://" + server;
        if (null != port) {
            serverBaseUrl = serverBaseUrl + ":" + port;
        }
        server_url = serverBaseUrl + "/api/datacenter/svm/svms";
        try {
            HttpURLConnection connection =
getAllTrustingHttpsURLConnection();
            if (connection == null) {
```

```

        System.err.println("FATAL: Failed to create HTTPS
connection to URL: " + server_url);
        System.exit(1);
    }
    System.out.println("Invoking API: " + server_url);
    connection.setRequestMethod("GET");
    connection.setRequestProperty("Accept", "application/" +
response_format);
    String authString = getAuthorizationString();
    connection.setRequestProperty("Authorization", "Basic " +
authString);
    if (connection.getResponseCode() != 200) {
        System.err.println("API Invocation Failed : HTTP error
code : " + connection.getResponseCode() + " : "
+ connection.getResponseMessage());
        System.exit(1);
    }
    BufferedReader br = new BufferedReader(new
InputStreamReader((connection.getInputStream())));
    String response;
    System.out.println("Response:");
    while ((response = br.readLine()) != null) {
        System.out.println(response);
    }
    connection.disconnect();
} catch (Exception e) {
    e.printStackTrace();
}
}

/* Print the usage of this sample code */ private static void
printUsage() {
    System.out.println("\nUsage:\n\tHelloApiServer <hostname> <user>
<password>\n");
    System.out.println("\nExamples:\n\tHelloApiServer localhost admin
mypassword");
    System.out.println("\tHelloApiServer 10.22.12.34:8320 admin
password");
    System.out.println("\tHelloApiServer 10.22.12.34 admin password
");
    System.out.println("\tHelloApiServer 10.22.12.34:8212 admin
password \n");
    System.out.println("\nNote:\n\t(1) When port number is not
provided, 443 is chosen by default.");
}
}

```

```

/* * Set the server, port, username and password * based on user
inputs. */ private static void setUserArguments(
    String[] args) {
    server = args[0];
    user = args[1];
    password = args[2];
    if (server.contains(":")) {
        String[] parts = server.split(":");
        server = parts[0];
        port = parts[1];
    }
}

/*
 * * Create a trust manager which accepts all certificates and * use
this trust
 * manager to initialize the SSL Context. * Create a
URLConnection for this
 * SSL Context and skip * server hostname verification during SSL
handshake. * *
 * Note: Trusting all certificates or skipping hostname verification *
is not
 * required for API Services to work. These are done here to * keep
this sample
 * REST Client code as simple as possible.
 */ private static HttpURLConnection
getAllTrustingHttpsURLConnection() {
    HttpURLConnection conn =
null;
    try {
        /* Creating a trust manager that does not
validate certificate chains */
        TrustManager[]
trustAllCertificatesManager = new
TrustManager[]{new
X509TrustManager() {
    public X509Certificate[] getAcceptedIssuers(){return null;}
    public void checkClientTrusted(X509Certificate[]
certs, String authType){}
    public void checkServerTrusted(X509Certificate[]
certs, String authType){}
}};
        /* Initialize the
SSLContext with the all-trusting trust manager */
        SSLContext sslContext = SSLContext.getInstance("TLS");
        sslContext.init(null, trustAllCertificatesManager, new
SecureRandom());
        HttpURLConnection.setDefaultSSLSocketFactory(sslContext.getSocketFactory(
));
        URL url = new URL(server_url);
        conn =
(HttpURLConnection) url.openConnection();
        /* Do not perform an
actual hostname verification during SSL Handshake.
Let all
hostname pass through as verified.*/
        conn.setHostnameVerifier(new HostnameVerifier() {
            public

```



```

boolean verify(String host, SSLSession session) {
return true; } catch (Exception e)
{ e.printStackTrace(); return conn; }

/*
 * * This forms the Base64 encoded string using the username and
password *
 * provided by the user. This is required for HTTP Basic
Authentication.
 */ private static String getAuthorizationString() {
String userPassword = user + ":" + password;
byte[] authEncodedBytes =
Base64.encodeBase64(userPassword.getBytes());
String authString = new String(authEncodedBytes);
return authString;
}
}

```

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, For the metrics and analytics information of the storage objects, see [Viewing performance metrics](#).

```
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.

APIs for storage objects in your data center

HTTP verb	Path	Description
GET	<code>/datacenter/cluster/clusters</code> <code>/datacenter/cluster/clusters/{key}</code>	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.

HTTP verb	Path	Description
GET	<pre>/datacenter/cluster/licensing/licenses /datacenter/cluster/licensing/licenses/{key}</pre>	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.
GET	<pre>/datacenter/cluster/nodes /datacenter/cluster/nodes/{key}</pre>	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	<pre>/datacenter/protocols/cifs/shares /datacenter/protocols/cifs/shares/{key}</pre>	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	<pre>/datacenter/protocols/nfs/export-policies /datacenter/protocols/nfs/export-policies/{key}</pre>	<p>You can use this method to view the details of the export policies for the supported NFS services.</p> <p>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”.</p>
GET	<pre>/datacenter/storage/aggregates /datacenter/storage/aggregates/{key}</pre>	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.

HTTP verb	Path	Description
GET	<pre>/datacenter/storage/luns /datacenter/storage/luns/{key}</pre>	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.
GET	<pre>/datacenter/storage/qos/policies /datacenter/storage/qos/policies/{key}</pre>	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.
GET	<pre>/datacenter/storage/qtrees /datacenter/storage/qtrees/{key}</pre>	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.
GET	<pre>/datacenter/storage/volumes /datacenter/storage/volumes/{key}</pre>	<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	<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"> • POST: While creating an igroup, you can designate the storage VM on which you want to assign access. • DELETE: You need to provide the igroup key as an input parameter to delete a particular igroup. If you have already assigned an igroup to a LUN, you cannot delete that igroup. • PATCH: You need to provide the igroup key as an input parameter to modify a particular igroup. You must also enter the property that you want to update, along with its value.
POST		
DELETE		
PATCH		

HTTP verb	Path	Description
GET	/datacenter/svm/svms	<p>You can use these methods to view, create, delete, and modify Storage Virtual Machines (storage VMs).</p> <ul style="list-style-type: none"> • 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. • DELETE: You need to provide the storage VM key to delete a particular storage VM. • PATCH: You need to provide the storage VM key to modify a particular storage VM. You also need to enter the properties that you want to update, along with their values.
POST	/datacenter/svm/svms/{key}	
DELETE		
PATCH		



Points to note:

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.

If you have enabled SLO-based workload provisioning on your environment, you cannot delete that storage VM on which storage workloads have been provisioned. When you delete a storage VM on which a CIFS or SMB server has been configured, this API also deletes the CIFS or SMB server, along with the local Active Directory configuration. However, the CIFS or SMB server name continues to be in the Active Directory configuration that you must delete manually from the Active Directory server.

APIs for network elements in your data center

The following APIs in the datacenter category retrieve information about the ports and network interfaces in your environment, specifically the FC ports, FC interfaces, ethernet ports, and IP interfaces.

HTTP verb	Path	Description
GET	<pre>/datacenter/network/ethernet/ports</pre> <pre>/datacenter/network/ethernet/ports/{key}</pre>	Retrieves information about all the ethernet ports in your data center environment. With a port key as an input parameter, you can view the information of that specific port. Information, such as cluster details, broadcast domain, port details, such as state, speed, and type, and whether the port is enabled, is retrieved.
GET	<pre>/datacenter/network/fc/interfaces</pre> <pre>/datacenter/network/fc/interfaces/{key}</pre>	You can use this method to view the details of the FC interfaces in your data center environment. With an interface key as an input parameter, you can view the information of that specific interface. Information, such as cluster details, home node details, and home port details, is retrieved.
GET	<pre>/datacenter/network/fc/ports</pre> <pre>/datacenter/network/fc/ports/{key}</pre>	Retrieves information about all the FC ports used in the nodes in your data center environment. With a port key as an input parameter, you can view the information of that specific port. Information, such as cluster details, port description, supported protocol, and the state of the port is retrieved.
GET	<pre>/datacenter/network/ip/interfaces</pre> <pre>/datacenter/network/ip/interfaces/{key}</pre>	You can use this method to view the details of the IP interfaces in your data center environment. With an interface key as an input parameter, you can view the information of that specific interface. Information, such as cluster details, IPspace details, home node details, whether failover is enabled, is retrieved.

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> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>The gateway APIs retrieve only those clusters supported by ONTAP 9.5 or later, and added to Unified Manager over HTTPS.</p> </div>

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		The <code>OPTIONS</code> method returns all the operations supported by an ONTAP REST API. For example, if an ONTAP API supports only the <code>GET</code> operation, running the <code>OPTIONS</code> method by using this gateway API returns <code>GET</code> as the response. This method is not supported on Swagger, but can be performed on other API tools.
PATCH		
OPTIONS (not available on Swagger)		The <code>OPTIONS</code> method determines whether a resource is available. This operation can be used to view the metadata about a resource in the HTTP response headers. This method is not supported on Swagger, but can be performed on other API tools.
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:

```
GEThttps://<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	<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"> • Whether the schedule is enabled or disabled • Frequency of the backup scheduled (daily or weekly) • Time of the backup • Maximum number of backup files that should be retained in the application <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	<p>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.</p>
GET	/admin/datasource-certificate	<p>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.</p>

HTTP verb	Path	Description
GET	/admin/datasources/clusters	<p>You can use the GET method to retrieve the details of the datasources (ONTAP clusters) managed by Unified Manager.</p> <p>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.</p> <p>For modifying and deleting a cluster managed as a datasource by Unified Manager, use the ONTAP cluster key.</p>
POST	/admin/datasources/clusters/{key}	
PATCH		
DELETE		

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	<p>You can use these methods to get the details of users or add new user to Unified Manager.</p> <p>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.</p>
POST	/security/users	

HTTP verb	Path	Description
GET	/security/users/{name}	The GET method enables you to retrieve all the details of a user, such as the name, email address, role, authorization type. The PATCH method enables you to update the details. The DELETE method enables you to remove the user.
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, storage VMs, FC interfaces, FC ports, Ethernet ports, and IP interfaces.

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
- FC interfaces
- FC ports
- Ethernet ports
- IP interfaces

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.

Metrics	Analytics
Performance metrics sample for a storage object based on the time interval parameter for retrieval.	The 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

Accumulated counters, such as minimum, maximum, 95th percentile, 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

`/datacenter/cluster/nodes/{key}/metrics` API retrieves the following details (among others) for a node :



The 95 percentile in the summary value indicates that 95% of the samples collected for the period have a counter value less than the value specified as the 95 percentile.

```
{
  "iops": {
    "local": {
      "other": 100.53,
      "read": 100.53,
      "total": 100.53,
      "write": 100.53
    },
    "other": 100.53,
    "read": 100.53,
    "total": 100.53,
    "write": 100.53
  },
  "latency": {
    "other": 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
  },
}
```

Analytics

A single aggregated value is displayed for all the objects of the same type. For example, the `/datacenter/cluster/nodes/analytics` API retrieves the following values (among others) for all the nodes:

```
{
  "iops": 1.7471,
  "latency": 60.0933,
  "throughput": 5548.4678,
  "utilization_percent":
4.8569,
  "period": 72,
  "performance_capacity": {
    "used_percent": 5.475,
    "available_iops_percent":
168350
  },
  "node": {
    "key": "37387241-8b57-11e9-
8974-
00a098e0219a:type=cluster_node,uui
d=95f94e8d-8b4e-11e9-8974-
00a098e0219a",
    "uuid": "95f94e8d-8b4e-
11e9-8974-00a098e0219a",
    "name": "ocum-infinity-01",
    "_links": {
      "self": {
        "href":
"/api/datacenter/cluster/nodes/373
87241-8b57-11e9-8974-
00a098e0219a:type=cluster_node,uui
d=95f94e8d-8b4e-11e9-8974-
00a098e0219a"
      }
    }
  },
  "cluster": {
    "key": "37387241-8b57-11e9-
8974-
00a098e0219a:type=cluster,uuid=373
87241-8b57-11e9-8974-
00a098e0219a",
    "uuid": "37387241-8b57-
11e9-8974-00a098e0219a",
    "name": "ocum-infinity",
    "_links": {
```

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"> • 1h: Metrics over the most recent hour sampled over 5 minutes. • 12h: Metrics over the most recent 12 hours sampled over 5 minutes. • 1d: Metrics over the most recent day sampled over 5 minutes. • 2d: Metrics over the most recent 2 days sampled over 5 minutes. • 3d: Metrics over the most recent 3 days sampled over 5 minutes. • 15d: Metrics over the most recent 15 days sampled over 1 hour. • 1w: Metrics over the most recent week sampled over 1 hour. • 1m: Metrics over the most recent month sampled over 1 hour. • 2m: Metrics over the most recent 2 months sampled over 1 hour. • 3m: Metrics over the most recent 3 months sampled over 1 hour. • 6m: Metrics over the most recent 6 months sampled over 1 hour. <p>Available values : 1h, 12h, 1d, 2d, 3d, 15d, 1w, 1m, 2m, 3m, 6m</p> <p>Default value : 1h</p>	<p>Above 72 hours. 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 the 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/analytcs /datacenter/storage/volumes/analytcs	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 <i>nvmf</i> , <i>fcp</i> , <i>iscsi</i> , and <i>nfs</i> , 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/network/ethernet/ports/{key}/metrics	Retrieves the performance metrics for a specific ethernet port specified by the input parameter of the port key. When an interval (time range) is provided from the supported range, the API returns the accumulated counters, such as minimum, maximum, and the average performance values over the period of time.
GET	/datacenter/network/ethernet/ports/analytics	Retrieves the high-level performance metrics for all the ethernet ports in your data center environment. Information, such as the cluster and node key and UUID, throughput, period of collection, and utilization percent for the ports is returned. You can filter the result by the available parameters, such as the port key, utilization percent, cluster and node name and UUID, and so forth.
GET	/datacenter/network/fc/interfaces/{key}/metrics	Retrieves the performance metrics for a specific network FC interface specified by the input parameter of the interface key. When an interval (time range) is provided from the supported range, the API returns the accumulated counters, such as minimum, maximum, and the average performance values over the period of time.

HTTP Verb	Path	Description
GET	/datacenter/network/fc/interfaces/analytics	Retrieves the high-level performance metrics for all the ethernet ports in your data center environment. Information, such as the cluster and FC interface key and UUID, throughput, IOPS, latency, and storage VM is returned. You can filter the result by the available parameters, such as the cluster and FC interface name and UUID, storage VM, throughput, and so forth.
GET	/datacenter/network/fc/ports/{key}/metrics	Retrieves the performance metrics for a specific FC port specified by the input parameter of the port key. When an interval (time range) is provided from the supported range, the API returns the accumulated counters, such as minimum, maximum, and the average performance values over the period of time.
GET	/datacenter/network/fc/ports/analytics	Retrieves the high-level performance metrics for all the FC ports in your data center environment. Information, such as the cluster and node key and UUID, throughput, period of collection, and utilization percent for the ports is returned. You can filter the result by the available parameters, such as the port key, utilization percent, cluster and node name and UUID, and so forth.
GET	/datacenter/network/ip/interfaces/{key}/metrics	Retrieves the performance metrics for a network IP interface as specified by the input parameter of the interface key. When an interval (time range) is provided from the supported range, the API returns information, such as the number of samples, accumulated counters, throughput, and the number of packets received and transmitted.

HTTP Verb	Path	Description
GET	/datacenter/network/ip/interfaces/analytics	Retrieves the high-level performance metrics for all the network IP interfaces in your data center environment. Information, such as the cluster and IP interface key and UUID, throughput, IOPS, and latency is returned. You can filter the result by the available parameters, such as the cluster and IP interface name and UUID, IOPS, latency, throughput, and so forth.

Viewing Jobs and system details

You can use the `jobs` API under the `management-server` category to view the execution details of asynchronous operations. The `system` API under the `management-server` category enables you to view the instance details in your Active IQ Unified Manager environment.

Viewing Jobs

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 you 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.
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 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	<code>/management-server/system</code>	No input parameter is required for running this API. The system details of the current Unified Manager instance are returned by default.

Managing events and alerts

The `events`, `alerts`, and `scripts` APIs under the `management-server` category enable you to manage the events, alerts, and the scripts associated with the alerts in your Active IQ Unified Manager environment.

Viewing and modifying events

Unified Manager receives the events that are generated on ONTAP for the clusters monitored and managed by Unified Manager. By using these APIs, you can view the events generated for your clusters, and resolve and update them.

By running the `GET` method for 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.

By running the `PATCH` operation for this API, you can enable the resolution workflow for the event. You can assign an event to yourself or another user and acknowledge the receipt of the event. On performing the steps on the resources to resolve the issue that triggered the event, you can use this API to mark the event as resolved.

For more information about events, see [Managing events](#)

Category	HTTP verb	Path	Description
management-server	GET	/management-server/events /management-server/events/{key}	When you run the Get ALL method, the response body consists of the event details of the all the events in your data center. When you retrieve the event details by a specific key, you can view the details of a specific event and run the next set of operations on the resources. The response body consists of the details of that event.
management-server	PATCH	management-server/events/{key}	Run this API to assign an event or change the state to acknowledged or resolved. You can also use this method to assign the event to yourself or another user. It is a synchronous operation.

Managing alerts

Events are automatically and continuously generated. Unified Manager generates an alert only when an event meets certain filter criteria. You can select the events for which alerts should be generated. By using the /management-server/alerts API, you can configure alerts to send notifications automatically when specific events or events of certain severity types occur.

For more information about alerts, see [Managing alerts](#)

Category	HTTP verb	Path	Description
management-server	GET	/management-server/alerts /management-server/alerts/{key}	Query all the existing alerts in your environment or a specific alert, by using the alert key. You can view the information about the alerts generated in your environment such as the alert description, action, email ID to which the notification is sent, event and severity.

Category	HTTP verb	Path	Description
management-server	POST	/management-server/alerts	This method enables you to add alerts for specific events. You must add the alert name, the physical or logical resource, or event on which the alert is applicable, whether the alert is enabled, and whether you are issuing SNMP traps. You can add additional details for which you want to generate the alert, such as the action, notification email ID, the script details, in case you are adding an alert script, and so forth.
management-server	PATCH and DELETE	management-server/events/{key}	You can use these methods to modify and delete specific alerts. You can modify different attributes, such as description, name, and enabling and disabling the alert. You can delete an alert when the alert is no longer required.



While selecting a resource for adding an alert, note that selecting a cluster as the resource does not automatically select the storage objects within that cluster. For example, if you create an alert for all the critical events for all clusters, you will receive alerts only for cluster critical events. You will not receive alerts for critical events on nodes, aggregates, and so forth.

Managing scripts

By using the `/management-server/scripts` API, you can also associate an alert with a script that is executed when an alert is triggered. You can use scripts to automatically modify or update multiple storage objects in Unified Manager. The script is associated with an alert. When an event triggers an alert, the script is executed. You can upload custom scripts and test their execution when an alert is generated. You can associate an alert with your script so that the script is executed when an alert is raised for an event in Unified Manager.

For more information about scripts, see [Managing scripts](#)

Category	HTTP verb	Path	Description
management-server	GET	/management-server/scripts	Use this API to query all the existing scripts in your environment. Use the standard filter and order by operations to view only specific scripts.
management-server	POST	/management-server/scripts	Use this API to add a description for the script and upload the script file associated with an alert.

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}

Common workflows for storage management

The common workflows provide client application developers with examples of how Active IQ Unified Manager APIs can be called by a client application to execute common storage management functions. This section contains some of these sample workflows.

The workflows describe some of the commonly used storage management use cases along with sample codes for you to use. Each of the tasks is described using a workflow process consisting of one or more API calls.

Understanding the API calls used in the workflows

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 used in the workflow samples 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 the 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 to the bottom of the page and 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 applies to as part of performing a call. The path string is appended to the core URL to form the complete URL identifying the resource.

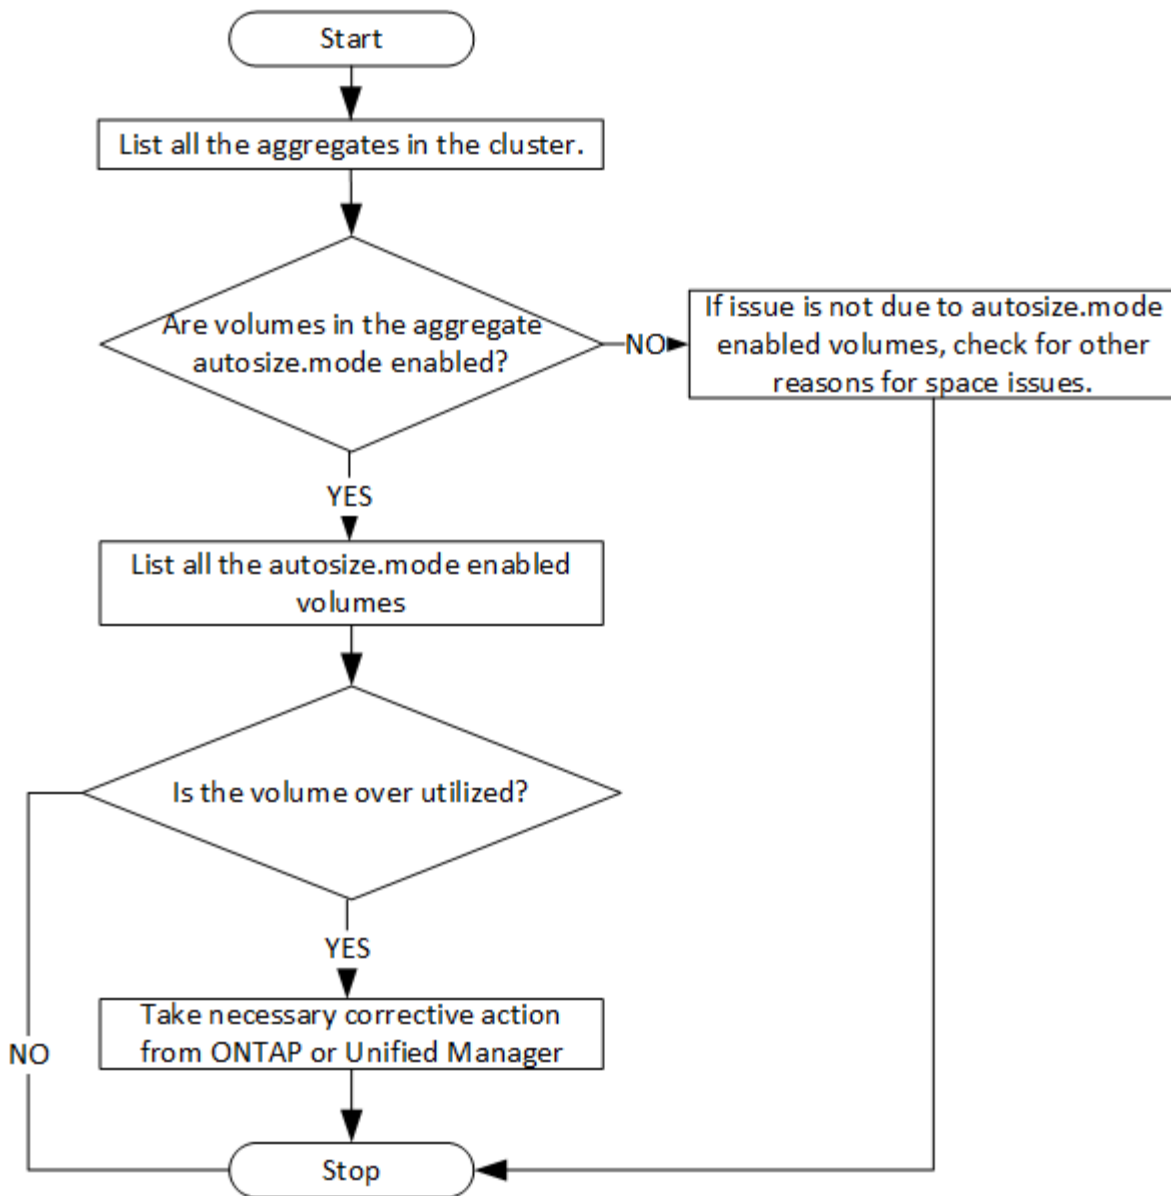
Determining space issues in aggregates

You can use the data center APIs in Active IQ Unified Manager to monitor the availability and utilization of space in your volumes. You can determine space issues in your volume and identify storage resources that are overutilized or underutilized.

The data center APIs for aggregates retrieve the relevant information about available and used space, and space saving efficiency settings. You can also filter the retrieved information based on specified attributes.

One method to determine any lack of space in your aggregates is to verify whether there are volumes in your environment with autosize-mode enabled. You should then identify which volumes are being over-utilized and perform any corrective actions.

The following flowchart illustrates the process of retrieving information about volumes with autosize-mode enabled:



This flow assumes that the clusters have already been created in ONTAP and added to Unified Manager.

1. Obtain the cluster key, unless you know the value:

Category	HTTP verb	Path
datacenter	GET	/datacenter/cluster/clusters

2. Using the cluster key as the filter parameter, query the aggregates on that cluster.

Category	HTTP verb	Path
datacenter	GET	/datacenter/storage/aggregates

3. From the response, analyze the space usage of the aggregates and determine which aggregates have space issues. For each aggregate with space issue, obtain the aggregate key from the same JSON output.
4. Using each aggregate key, filter all the volumes that have the value for the `autosize.mode` parameter as `grow`.

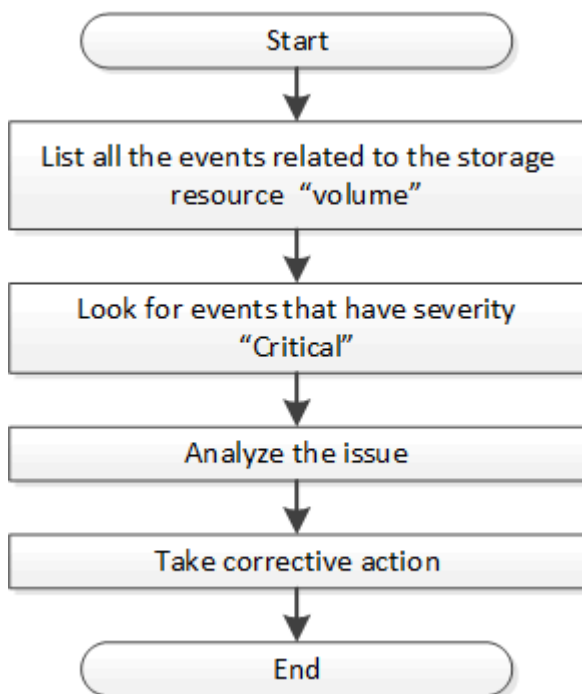
Category	HTTP verb	Path
datacenter	GET	/datacenter/storage/volumes

5. Analyze which volumes are being over-utilized.
6. Perform any necessary corrective action, such as moving the volume across aggregates, to address the space issues in your volume. You can perform these actions from ONTAP or Unified Manager web UI.

Determining issues in storage objects using events

When a storage object in your data center crosses a threshold, you get a notification about that event. Using this notification, you can analyze the issue and take corrective action by using the `events` APIs.

This workflow takes the example of a volume as the resource object. You can use the `events` APIs to retrieve the list of events related to a volume, analyze the critical issues for that volume, and then take corrective actions to rectify the issue.



Follow these steps to determine the issues in your volume before taking remedial steps.

Steps

1. Analyze the critical Active IQ Unified Manager events notifications for the volumes in your data center.
2. Query all the events for the volumes by using the following parameters in the `/management-server/events`

API: "resource_type": "volume" "severity": "critical"

Category	HTTP verb	Path
management-server	GET	/management-server/events

3. View the output and analyze the issues in the specific volumes.
4. Perform the necessary actions by using the Unified Manager REST APIs or web UI to resolve the issues.

Troubleshooting ONTAP volumes by using gateway APIs

The gateway APIs act as a gateway to invoke ONTAP APIs to query information about your ONTAP storage objects and take remedial measures to address the reported issues.

This workflow takes up a sample use case in which an event is raised when an ONTAP volume almost reaches its capacity. The workflow also demonstrates how to address this issue by invoking a combination of Active IQ Unified Manager and ONTAP REST APIs.

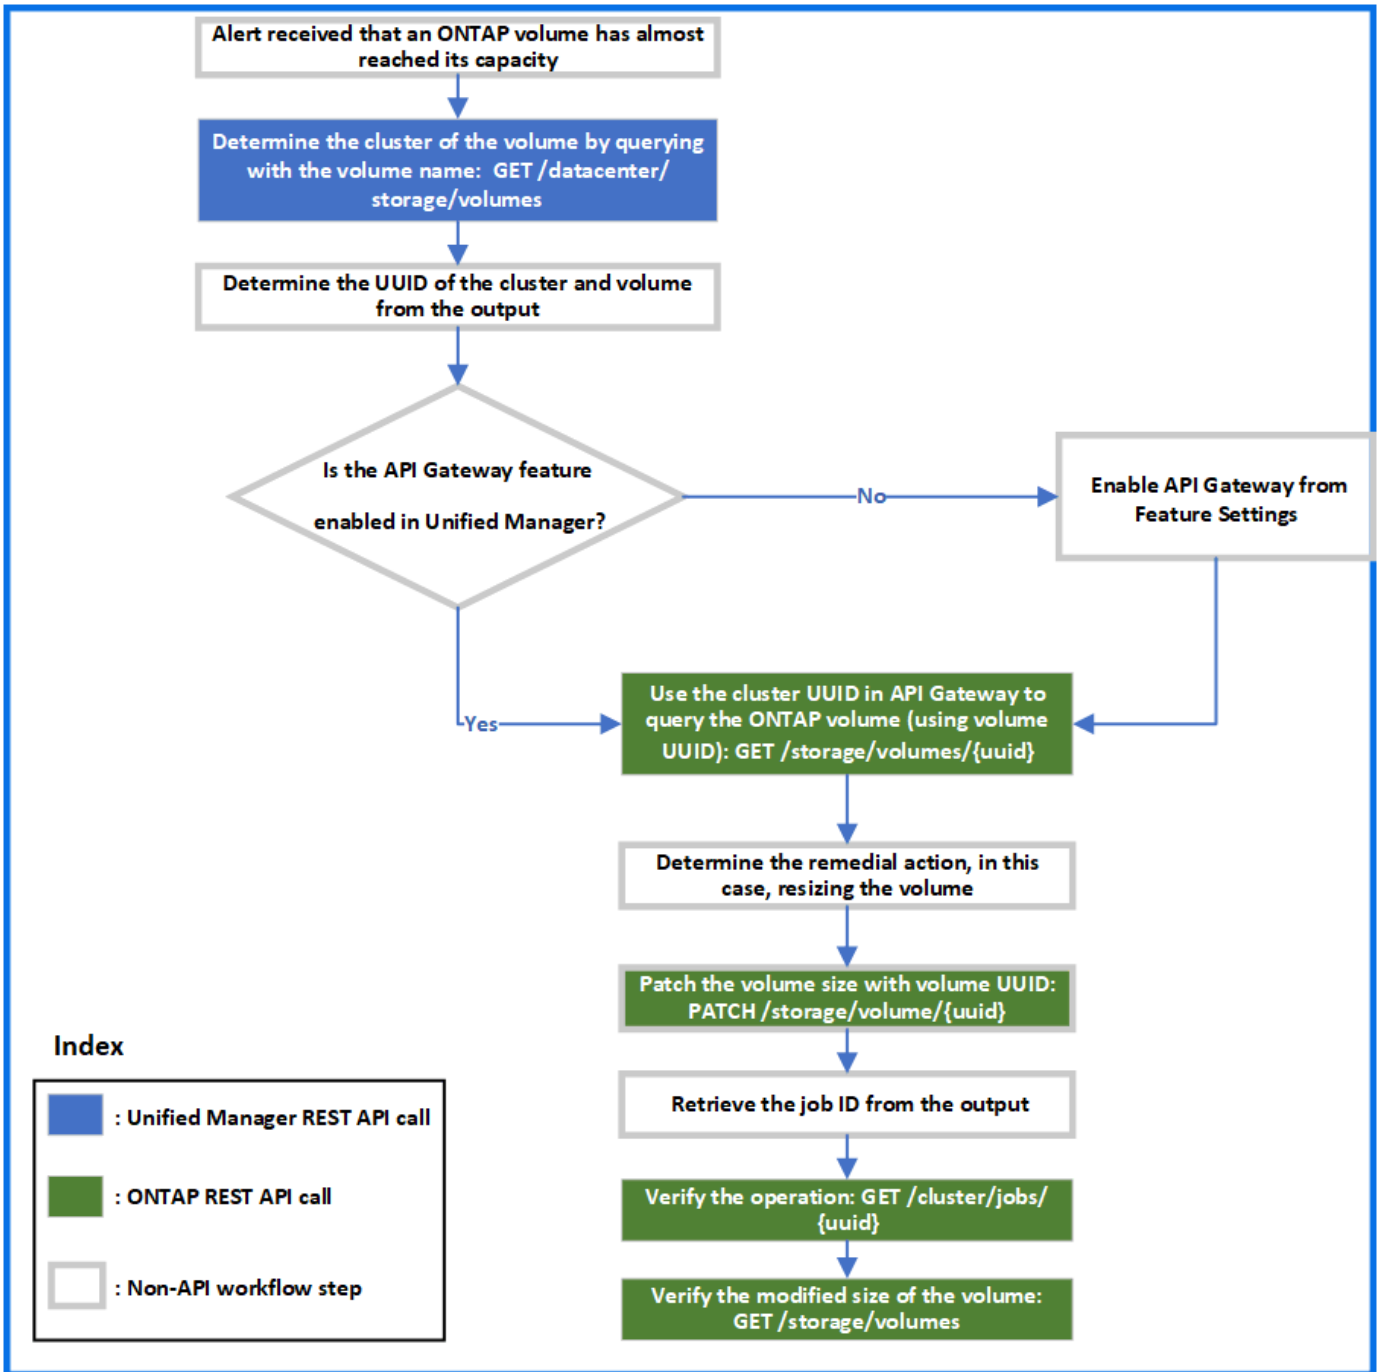
Before running the workflow steps, ensure the following:

- You are aware of the gateway APIs and how they are used. For information, see the “Gateway APIs” section.

[Accessing ONTAP APIs through proxy access](#)

- You are aware of the usage of ONTAP REST APIs. For information about using ONTAP REST APIs, see [ONTAP Automation documentation](#).
- You are an Application Administrator.
- The cluster on which you want to run the REST API operations is supported by ONTAP 9.5 or later, and the cluster is added to Unified Manager over HTTPS.

The following diagram illustrates each step in the workflow for troubleshooting the issue of ONTAP volume capacity use.



The workflow covers the invocation points of both the Unified Manager and ONTAP REST APIs.

1. Note the volume name from the event notifying the volume capacity utilization.
2. By using the volume name as the value in the name parameter, query the volume by running the following Unified Manager API.

Category	HTTP verb	Path
datacenter	GET	/datacenter/storage/volumes

3. Retrieve the cluster UUID and volume UUID from the output.

- On the Unified Manager web UI, navigate to **General > Feature Settings > API Gateway** to verify whether the API Gateway feature is enabled. Unless it is enabled, the APIs under the gateway category are not available for you to invoke. Enable the feature if it is disabled.
- Use the cluster UUID to run the ONTAP API `/storage/volumes/{uuid}` through API gateway. The query returns the volume details when the volume UUID is passed as the API parameter.

For running the ONTAP APIs through API gateway, the Unified Manager credentials are passed internally for authentication, and you do not need to run an additional authentication step for individual cluster access.

Category	HTTP verb	Path
Unified Manager: gateway	GET	Gateway API: <code>/gateways/{uuid}/{path}</code>
ONTAP: storage		ONTAP API: <code>/storage/volumes/{uuid}</code>



In `/gateways/{uuid}/{path}`, the value for `{uuid}` must be replaced with the cluster UUID on which the REST operation is to be performed. `{path}` must be replaced by the ONTAP REST URL `/storage/volumes/{uuid}`.

The appended URL is: `/gateways/{cluster_uuid}/storage/volumes/{volume_uuid}`

On running the GET operation, the generated URL is:

```
GEThttps://<hostname>/api/gateways/<cluster_UUID>/storage/volumes/{volume_uuid}
```

Sample cURL command

```
curl -X GET "https://<hostname>/api/gateways/1cd8a442-86d1-11e0-ae1c-9876567890123/storage/volumes/028baa66-41bd-11e9-81d5-00a0986138f7"
-H "accept: application/hal+json" -H "Authorization: Basic
<Base64EncodedCredentials>"
```

- From the output, determine the size, usage, and remedial measure to be taken. In this workflow, the remedial measure taken is resizing the volume.
- Use the cluster UUID and run the following ONTAP API through the API gateway to resize the volume. For information about the input parameters for the gateway and ONTAP APIs, see step 5.

Category	HTTP verb	Path
Unified Manager: gateway	PATCH	Gateway API: <code>/gateways/{uuid}/{path}</code>
ONTAP: storage		ONTAP API: <code>/storage/volumes/{uuid}</code>



Along with the cluster UUID and volume UUID, you must enter a value for the size parameter for resizing the volume. Ensure to enter the value *in bytes*. For example, if you want to increase the size of a volume from 100 GB to 120 GB, enter the value for parameter size at the end of the query: `-d {"size": 128849018880}"`

Sample cURL command

```
curl -X PATCH "https://<hostname>/api/gateways/1cd8a442-86d1-11e0-ae1c-9876567890123/storage/volumes/028baa66-41bd-11e9-81d5-00a0986138f7" -H "accept: application/hal+json" -H "Authorization: Basic <Base64EncodedCredentials>" -d {"size": 128849018880}"
```

The JSON output returns a Job UUID.

- Verify whether the job ran successfully by using the Job UUID. Use the cluster UUID and Job UUID to run the following ONTAP API through the API gateway. For information about the input parameters for the gateway and ONTAP APIs, see step 5.

Category	HTTP verb	Path
Unified Manager: gateway	GET	Gateway API: /gateways/{uuid}/{path}
ONTAP: cluster		ONTAP API: /cluster/jobs/{uuid}

The HTTP codes returned are the same as the ONTAP REST API HTTP status codes.

- Run the following ONTAP API to query the details of the resized volume. For information about the input parameters for the gateway and ONTAP APIs, see step 5.

Category	HTTP verb	Path
Unified Manager: gateway	GET	Gateway API: /gateways/{uuid}/{path}
ONTAP: storage		ONTAP API: /storage/volumes/{uuid}

The output displays an increased volume size of 120 GB.

Workflows for workload management

Using Active IQ Unified Manager, you can provision and modify storage workloads (LUNs, NFS file shares, and CIFS shares). Provisioning consists of multiple steps, from the creation of the Storage Virtual Machine (SVM) to applying Performance Service Level and Storage Efficiency Policies on the storage workloads. Modifying workloads consist of

the steps for modifying specific parameters and enabling additional features on them.

The following workflows are described:

- Workflow for provisioning Storage Virtual Machines (SVMs) on Unified Manager.



this workflow is required to be performed before provisioning LUNs or file shares on Unified Manager.

- Provisioning file shares.
- Provisioning LUNs.
- Modifying LUNs and file shares (by using the example for updating the Performance Service Level parameter for the storage workloads).
- Modifying an NFS file share to support CIFS protocol
- Modifying workloads to upgrade QoS to AQoS



For each provisioning workflow (LUN and file shares), ensure you must have completed the workflow for verifying the SVMs on the clusters.

You must also read the recommendations and limitations before using each API in the workflows. The relevant details of the APIs are available in their individual sections listed in the related concepts and references.

Verifying SVMs on clusters

Before provisioning file shares or LUNs, you must verify whether the clusters have Storage Virtual Machines (SVMs) created on them.



The workflow assumes that ONTAP clusters to have been added to Unified Manager, and the cluster key has been obtained. Clusters should have the required licenses for provisioning LUNs and file shares on them.

1. Verify whether the cluster has an SVM created.

Category	HTTP verb	Path
datacenter	GET	/datacenter/svm/svms /datacenter/svm/svms/{key }

Sample cURL

```
curl -X GET "https://<hostname>/api/datacenter/svm/svms" -H "accept: application/json" -H "Authorization: Basic <Base64EncodedCredentials>"
```

2. If the SVM key is not returned, then create the SVM. For creating the SVMs, you require the cluster key on which you provision the SVM. You also need to specify the SVM name. Follow these steps.

Category	HTTP verb	Path
datacenter	GET	/datacenter/cluster/clusters /datacenter/cluster/clusters/{key}

Get the cluster key.

Sample cURL

```
curl -X GET "https://<hostname>/api/datacenter/cluster/clusters" -H
"accept: application/json" -H "Authorization: Basic
<Base64EncodedCredentials>"
```

- From the output, get the cluster key, and then use it as an input for creating the SVM.



While creating the SVM, ensure that it supports all the protocols required for provisioning LUNs and file shares on them, for example, CIFS, NFS, FCP, and iSCSI. The provisioning workflows might fail if the SVM does not support the required services. It is recommended that the services for the respective types of workloads are also enabled on the SVM.

Category	HTTP verb	Path
datacenter	POST	/datacenter/svm/svms

Sample cURL

Enter the SVM object details as input parameters.

```

curl -X POST "https://<hostname>/api/datacenter/svm/svms" -H "accept:
application/json" -H "Content-Type: application/json" -H "Authorization:
Basic <Base64EncodedCredentials>" "{ \"aggregates\": [ { \"_links\": {},
\"key\": \"1cd8a442-86d1,type=objecttype,uuid=1cd8a442-86d1-11e0-ae1c-
9876567890123\",
\"name\": \"cluster2\", \"uuid\": \"02c9e252-41be-11e9-81d5-
00a0986138f7\" } ],
\"cifs\": { \"ad_domain\": { \"fqdn\": \"string\", \"password\":
\"string\",
\"user\": \"string\" }, \"enabled\": true, \"name\": \"CIFS1\" },
\"cluster\": { \"key\": \"1cd8a442-86d1-11e0-ae1c-
123478563412,type=object type,uuid=1cd8a442-86d1-11e0-ae1c-
9876567890123\" },
\"dns\": { \"domains\": [ \"example.com\", \"example2.example3.com\" ],
\"servers\": [ \"10.224.65.20\", \"2001:db08:a0b:12f0::1\" ] },
\"fcg\": { \"enabled\": true }, \"ip_interface\": [ { \"enabled\": true,
\"ip\": { \"address\": \"10.10.10.7\", \"netmask\": \"24\" },
\"location\": { \"home_node\": { \"name\": \"node1\" } }, \"name\":
\"dataLif1\" } ], \"ipspace\": { \"name\": \"exchange\" },
\"iscsi\": { \"enabled\": true }, \"language\": \"c.utf_8\",
\"ldap\": { \"ad_domain\": \"string\", \"base_dn\": \"string\",
\"bind_dn\": \"string\", \"enabled\": true, \"servers\": [ \"string\" ]
},
\"name\": \"svm1\", \"nfs\": { \"enabled\": true },
\"nis\": { \"domain\": \"string\", \"enabled\": true,
\"servers\": [ \"string\" ] }, \"nvme\": { \"enabled\": true },
\"routes\": [ { \"destination\": { \"address\": \"10.10.10.7\",
\"netmask\": \"24\" }, \"gateway\": \"string\" } ],
\"snapshot_policy\": { \"name\": \"default\" },
\"state\": \"running\", \"subtype\": \"default\"}"

```

The JSON output displays a Job object key that you can use to verify the SVM that you created.

4. Verify the SVM creation by using the job object key for query. If the SVM is created successfully, the SVM key is returned in the response.


Category	HTTP verb	Path
management-server	GET	/management-server/jobs/{key}

Provisioning CIFS and NFS file shares

You can provision CIFS shares and NFS file shares on your Storage Virtual Machines (SVMs) by using the provisioning APIs provided as a part of Active IQ Unified Manager.

This provisioning workflow details the steps for retrieving the keys of the SVMs, Performance Service Levels, and Storage Efficiency Policies before creating the file shares.

The following diagram illustrates each step in a file share provisioning workflow. It includes provisioning both CIFS shares and NFS file shares.



Ensure the following:



- ONTAP clusters have been added to Unified Manager, and the cluster key has been obtained.
- SVMs have been created on the clusters.
- The SVMs support CIFS and NFS services. Provisioning file shares might fail if the SVMs do not support the required services.
- The FCP port is online for port provisioning.

1. Determine whether Data LIFs or access endpoints are available on the SVM on which you want to create the CIFS share. Get the list of available access endpoints on the SVM:

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

Sample cURL

```
curl -X GET "https://<hostname>/api/storage-provider/access-endpoints?resource.key=7d5a59b3-953a-11e8-8857-00a098dcc959" -H "accept: application/json" -H "Authorization: Basic <Base64EncodedCredentials>"
```

2. If your access endpoint is available on the list, obtain the access endpoint key, else create the access endpoint.



Ensure that you create access endpoints that have the CIFS protocol enabled on them. Provisioning CIFS shares fails unless you have created an access endpoint with the CIFS protocol enabled on it.

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

Sample cURL

You must enter the details of the access endpoint that you want to create, as input parameters.

```
curl -X POST "https://<hostname>/api/storage-provider/access-endpoints"
-H "accept: application/json" -H "Content-Type: application/json" -H
"Authorization: Basic <Base64EncodedCredentials>"
{ \"data_protocols\": \"nfs\",
\"fileshare\": { \"key\": \"cbd1757b-0580-11e8-bd9d-
00a098d39e12:type=volume,uuid=f3063d27-2c71-44e5-9a69-a3927c19c8fc\" },
\"gateway\": \"10.132.72.12\",
\"ip\": { \"address\": \"10.162.83.26\",
\"ha_address\": \"10.142.83.26\",
\"netmask\": \"255.255.0.0\" },
\"lun\": { \"key\": \"cbd1757b-0580-11e8-bd9d-
00a098d39e12:type=lun,uuid=d208cc7d-80a3-4755-93d4-5db2c38f55a6\" },
\"mtu\": 15000, \"name\": \"aep1\",
\"svm\": { \"key\": \"cbd1757b-0580-11e8-bd9d-
00a178d39e12:type=vserver,uuid=1d1c3198-fc57-11e8-99ca-00a098d38e12\" },
\"vlan\": 10}"
```

The JSON output displays a Job object key that you can use to verify the access endpoint that you created.

3. Verify the access endpoint:

Category	HTTP verb	Path
management-server	GET	/management-server/jobs/{key}

4. Determine whether you have to create a CIFS share or an NFS file share. For creating CIFS shares, follow these substeps:

- a. Determine whether the CIFS server is configured on your SVM, that is determine whether an Active Directory mapping is created on the SVM.

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

- b. If the Active Directory mapping is created, take the key, else create the Active Directory mapping on the SVM.

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

Sample cURL

You must enter the details for creating the Active Directory mapping, as the input parameters.

```
curl -X POST "https://<hostname>/api/storage-provider/active-
directories-mappings" -H "accept: application/json" -H "Content-Type:
application/json" -H "Authorization: Basic
<Base64EncodedCredentials>"
{ \"_links\": {},
  \"dns\": \"10.000.000.000\",
  \"domain\": \"example.com\",
  \"password\": \"string\",
  \"svm\": { \"key\": \"9f4ddea-e395-11e9-b660-
005056a71be9:type=vserver,uuid=191a554a-f0ce-11e9-b660-005056a71be9\"
},
  \"username\": \"string\"}
```

This is a synchronous call and you can verify the creation of the Active Directory mapping in the output. In case of an error, the error message is displayed for you to troubleshoot and rerun the request.

5. Obtain the SVM key for the SVM on which you want to create the CIFS share or the NFS file share, as described in the *Verifying SVMs on clusters* workflow topic.
6. Obtain the key for the Performance Service Level by running the following API and retrieving the key from the response.

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



You can retrieve the details of the system-defined Performance Service Levels by setting the `system_defined` input parameter to `true`. From the output, obtain the key of the Performance Service Level that you want to apply on the file share.

7. Optionally, obtain the Storage Efficiency Policy key for the Storage Efficiency Policy that you want to apply on the file share by running the following API and retrieving the key from the response.

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

8. Create the file share. You can create a file share that supports both CIFS and NFS by specifying the access control list and export policy. The following substeps provide information if you want to create a file share for supporting only one of the protocols on the volume. You can also update an NFS file share to include the access control list after you have created the NFS share. For information, see the *Modifying*

storage workloads topic.

- a. For creating only a CIFS share, gather the information about access control list (ACL). For creating the CIFS share, provide valid values for the following input parameters. For each user group that you assign, an ACL is created when a CIFS/SMB share is provisioned. Based on the values you enter for ACL and Active Directory mapping, the access control and mapping are determined for the CIFS share when it is created.

A cURL command with sample values

```
{
  "access_control": {
    "acl": [
      {
        "permission": "read",
        "user_or_group": "everyone"
      }
    ],
    "active_directory_mapping": {
      "key": "3b648c1b-d965-03b7-20da-61b791a6263c"
    }
  },
}
```

- b. For creating only an NFS file share, gather the information about the export policy. For creating the NFS file share, provide valid values for the following input parameters. Based on your values, the export policy is attached with the NFS file share when it is created.



While provisioning the NFS share, you can either create an export policy by providing all the required values or provide the export policy key and reuse an existing export policy. If you want to reuse an export policy for the storage VM, you need to add the export policy key. Unless you know the key, you can retrieve the export policy key by using the `/datacenter/protocols/nfs/export-policies` API. For creating a new policy, you must enter the rules as displayed in the following sample. For the entered rules, the API tries to search for an existing export policy by matching the host, storage VM, and rules. If there is an existing export policy, it is used. Otherwise a new export policy is created.

A cURL command with sample values

```

"export_policy": {
  "key": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=export_policy,uuid=1460288880641",
  "name_tag": "ExportPolicyNameTag",
  "rules": [
    {
      "clients": [
        {
          "match": "0.0.0.0/0"
        }
      ]
    }
  ]
}

```

After configuring access control list and export policy, provide the valid values for the mandatory input parameters for both CIFS and NFS file shares:



Storage Efficiency Policy is an optional parameter for creating file shares.

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

The JSON output displays a Job object key that you can use to verify the file share that you created. . Verify the file share creation by using the Job object key returned in querying the job:

Category	HTTP verb	Path
management-server	GET	/management-server/jobs/{key}

At the end of the response, you see the key of the file share created.

```

],
"job_results": [
  {
    "name": "fileshareKey",
    "value": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=volume,uuid=e581c23a-1037-11ea-ac5a-00a098dcc6b6"
  }
],
"_links": {
  "self": {
    "href": "/api/management-server/jobs/06a6148bf9e862df:-
2611856e:16e8d47e722:-7f87"
  }
}
}

```

1. Verify the creation of the file share by running the following API with the returned key:

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

Sample JSON output

You can see that the POST method of /storage-provider/file-shares internally invokes all the APIs required for each of the functions and creates the object. For example, it invokes the /storage-provider/performance-service-levels/ API for assigning the Performance Service Level on the file share.

```

{
  "key": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=volume,uuid=e581c23a-1037-11ea-ac5a-00a098dcc6b6",
  "name": "FileShare_377",
  "cluster": {
    "uuid": "7d5a59b3-953a-11e8-8857-00a098dcc959",
    "key": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=cluster,uuid=7d5a59b3-953a-11e8-8857-00a098dcc959",
    "name": "AFFA300-206-68-70-72-74",
    "_links": {
      "self": {
        "href": "/api/datacenter/cluster/clusters/7d5a59b3-953a-
11e8-8857-00a098dcc959:type=cluster,uuid=7d5a59b3-953a-11e8-8857-
00a098dcc959"
      }
    }
  }
}

```

```


    },
    "svm": {
      "uuid": "b106d7b1-51e9-11e9-8857-00a098dcc959",
      "key": "7d5a59b3-953a-11e8-8857-00a098dcc959:type=vserver,uuid=b106d7b1-51e9-11e9-8857-00a098dcc959",
      "name": "RRT_ritu_vs1",
      "_links": {
        "self": {
          "href": "/api/datacenter/svm/svms/7d5a59b3-953a-11e8-8857-00a098dcc959:type=vserver,uuid=b106d7b1-51e9-11e9-8857-00a098dcc959"
        }
      }
    },
    "assigned_performance_service_level": {
      "key": "1251e51b-069f-11ea-980d-fa163e82bbf2",
      "name": "Value",
      "peak_iops": 75,
      "expected_iops": 75,
      "_links": {
        "self": {
          "href": "/api/storage-provider/performance-service-levels/1251e51b-069f-11ea-980d-fa163e82bbf2"
        }
      }
    },
    "recommended_performance_service_level": {
      "key": null,
      "name": "Idle",
      "peak_iops": null,
      "expected_iops": null,
      "_links": {}
    },
    "space": {
      "size": 104857600
    },
    "assigned_storage_efficiency_policy": {
      "key": null,
      "name": "Unassigned",
      "_links": {}
    },
    "access_control": {
      "acl": [
        {
          "user_or_group": "everyone",

```

```

        "permission": "read"
    }
],
"export_policy": {
    "id": 1460288880641,
    "key": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=export_policy,uuid=1460288880641",
    "name": "default",
    "rules": [
        {
            "anonymous_user": "65534",
            "clients": [
                {
                    "match": "0.0.0.0/0"
                }
            ],
            "index": 1,
            "protocols": [
                "nfs3",
                "nfs4"
            ],
            "ro_rule": [
                "sys"
            ],
            "rw_rule": [
                "sys"
            ],
            "superuser": [
                "none"
            ]
        },
        {
            "anonymous_user": "65534",
            "clients": [
                {
                    "match": "0.0.0.0/0"
                }
            ],
            "index": 2,
            "protocols": [
                "cifs"
            ],
            "ro_rule": [
                "ntlm"
            ],
            "rw_rule": [

```

This workflow assumes that the ONTAP clusters have been added to Unified Manager, and the cluster key has been obtained. The workflow also assumes that the SVMs have already been created on the clusters.

1. Obtain the SVM key for the SVM on which you want to create the LUN, as described in the *Verifying SVMs on clusters* workflow topic.
2. Obtain the key for the Performance Service Level by running the following API and retrieving the key from the response.

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



You can retrieve the details of the system-defined Performance Service Levels by setting the `system_defined` input parameter to `true`. From the output, obtain the key of the Performance Service Level that you want to apply on the LUN.

- Optionally, obtain the Storage Efficiency Policy key for the Storage Efficiency Policy that you want to apply on the LUN by running the following API and retrieving the key from the response.

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

- Determine if initiator groups (igroups) have been created to grant access to the LUN target that you want to create.

Category	HTTP verb	Path
datacenter	GET	/datacenter/protocols/san/igroups /datacenter/protocols/san/igroups/{key}

You must enter the parameter value for indicating the SVM for which the igroup has authorized access. Additionally, if you want to query a particular igroup, enter the group name (`key`) as an input parameter.

- In the output, if you can find the igroup that you want to grant access to, obtain the key. Otherwise create the igroup.

Category	HTTP verb	Path
datacenter	POST	/datacenter/protocols/san/igroups

You must enter the details of the igroup that you want to create, as the input parameters. This is a synchronous call and you can verify the igroup creation in the output. In case of an error, a message is displayed for you to troubleshoot and rerun the API.

- Create the LUN.

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

For creating the LUN, ensure that you have added the retrieved values as mandatory input parameters.



Storage Efficiency Policy is an optional parameter for creating LUNs.

Sample cURL

You must enter all the details of the LUN that you want to create, as the input parameters.

Wombat: Extract # 1

The JSON output displays a Job object key that you can use to verify the LUN that you created.

- Verify the LUN creation by using the Job object key returned in querying the Job:

Category	HTTP verb	Path
management-server	GET	/management-server/jobs/{key}

At the end of the response, you see the key of the LUN created.

Wombat: Extract # 2

- Verify the creation of the LUN by running the following API with the returned key:

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

Sample JSON output

You can see that the POST method of /storage-provider/luns internally invokes all the APIs required for each of the functions and creates the object. For example, it invokes the /storage-provider/performance-service-levels/ API for assigning the Performance Service Level on the LUN.

Wombat: Extract # 3

Troubleshooting steps for failure in LUN creation or mapping

On completing this workflow, you might still see a failure in your LUN creation. Even if the LUN is created successfully, the LUN mapping with the igroup might fail due to an unavailability of a SAN LIF or access endpoint on the node on which you create the LUN. In case of a failure, you can see the following message:

The nodes <node_name> and <partner_node_name> have no LIFs configured with the iSCSI or FCP protocol for Vserver <server_name>. Use the access-endpoints API to create a LIF for the LUN.

Follow these troubleshooting steps to work around this failure.

1. Create an access endpoint supporting iSCSI/FCP protocol on the SVM on which you tried creating the LUN.

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

Sample cURL

You must enter the details of the access endpoint that you want to create, as the input parameters.



Ensure that in the input parameter you have added the address to indicate the home node of the LUN and the ha_address to indicate the partner node of the home node. When you run this operation, it creates access endpoints on both the home node and the partner node.

Wombat: Extract # 4

2. Query the job with the Job object key returned in the JSON output to verify that it has run successfully to add the access endpoints on the SVM and that the iSCSI/FCP services have been enabled on the SVM.

Category	HTTP verb	Path
management-server	GET	/management-server/jobs/{key}

Sample JSON output

At the end of the output, you can see the key of the access endpoints created. In the following output, the "name": "accessEndpointKey" value indicates the access endpoint created on the home node of the LUN, for which the key is 9c964258-14ef-11ea-95e2-00a098e32c28. The "name": "accessEndpointHAKey" value indicates the access endpoint created on the partner node of the home node, for which the key is 9d347006-14ef-11ea-8760-00a098e3215f.

Wombat: Extract # 5

3. Modify the LUN to update the igroup mapping. For more information about workflow modification, see "Modifying storage workloads".

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

In the input, specify the igroup key with which you want to update the LUN mapping, along with the LUN key.

Sample cURL

Wombat: Extract # 6

The JSON output displays a Job object key that you can use to verify whether the mapping is successful.

4. Verify the LUN mapping by querying with the LUN key.

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

Sample JSON output

In the output you can see the LUN has been successfully mapped with the igroup (key d19ec2fa-fec7-11e8-b23d-00a098e32c28) with which it was initially provisioned.

Wombat: Extract # 7

Modifying storage workloads

Modifying storage workloads consists of updating LUNs or file shares with missing parameters, or changing the existing parameters.

This workflow takes the example of updating Performance Service Levels for LUNs and file shares.



The workflow assumes that the LUN or file share has been provisioned with Performance Service Levels.

Modifying file shares

While modifying a file share, you can update the following parameters:

- Capacity or size.
- Online or offline setting.
- Storage Efficiency Policy.
- Performance Service Level.
- Access control list (ACL) settings.
- Export policy settings. You can also delete export policy parameters and revert the default (empty) export policy rules on the file share.



During a single API run, you can update only one parameter.

This procedure describes adding a Performance Service Level to a file share. You can use the same procedure for updating any other file share property.

1. Obtain the CIFS share or NFS file share key of the file share that you want to update. This API queries all the file shares on your data center. Skip this step if you already know the file share key.

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

2. View the details of the file share by running the following API with the file share key that you obtained.

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

View the details of the file share in the output.

```
"assigned_performance_service_level": {
  "key": null,
  "name": "Unassigned",
  "peak_iops": null,
  "expected_iops": null,
  "_links": {}
},
```

3. Obtain the key for the Performance Service Level that you want to assign on this file share. Currently no policy is assigned to it.

Category	HTTP verb	Path
Performance Service Levels	GET	/storage-provider/performance-service-levels



You can retrieve the details of the system-defined Performance Service Levels by setting the `system_defined` input parameter to `true`. From the output, obtain the key of the Performance Service Level that you want to apply to the file share.

4. Apply the Performance Service Level on the file share.

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

In the input, you must specify only the parameter that you want to update, along with the file share key. In this case, it is the key of the Performance Service Level.

Sample cURL

```
curl -X POST "https://<hostname>/api/storage-provider/file-shares" -H
"accept: application/json" -H "Authorization: Basic
<Base64EncodedCredentials>" -d
"{
  \"performance_service_level\": { \"key\": \"1251e51b-069f-11ea-980d-
fa163e82bbf2\" },
}"
```

The JSON output displays a Job object that you can use to verify the whether the access endpoints on the home and partner nodes have been created successfully.

5. Verify whether the Performance Service Level has been added to the file share by using the Job object key displayed in your output.

Category	HTTP verb	Path
Management Server	GET	/management-server/jobs/{key}

If you query by the ID of the Job object, you see whether the file share is updated successfully. In case of a failure, troubleshoot the failure and run the API again. On successful creation, query the file share to see the modified object:

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

View the details of the file share in the output.

```
"assigned_performance_service_level": {
  "key": "1251e51b-069f-11ea-980d-fa163e82bbf2",
  "name": "Value",
  "peak_iops": 75,
  "expected_iops": 75,
  "_links": {
    "self": {
      "href": "/api/storage-provider/performance-service-
levels/1251e51b-069f-11ea-980d-fa163e82bbf2"
    }
  }
}
```

Updating LUNs

While updating a LUN, you can modify the following parameters:

- Capacity or size
- Online or offline setting
- Storage Efficiency Policy
- Performance Service Level
- LUN map



During a single API run, you can update only one parameter.

This procedure describes adding a Performance Service Level to a LUN. You can use the same procedure for updating any other LUN property.

1. Obtain the LUN key of the LUN that you want to update. This API returns details of all the LUNS in your data center. Skip this step if you already know the LUN key.

Category	HTTP verb	Path
Storage Provider	GET	/storage-provider/luns

2. View the details of the LUN by running the following API with the LUN key that you obtained.

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

View the details of the LUN in the output. You can see that there is no Performance Service Level assigned to this LUN.

Sample JSON output

```
"assigned_performance_service_level": {
  "key": null,
  "name": "Unassigned",
  "peak_iops": null,
  "expected_iops": null,
  "_links": {}
},
```

3. Obtain the key for the Performance Service Level that you want to assign to the LUN.

Category	HTTP verb	Path
Performance Service Levels	GET	/storage-provider/performance-service-levels



You can retrieve the details of the system-defined Performance Service Levels by setting the `system_defined` input parameter to `true`. From the output, obtain the key of the Performance Service Level that you want to apply on the LUN.

4. Apply the Performance Service Level on the LUN.

Category	HTTP verb	Path
Storage Provider	PATCH	/storage-provider/lun/{key}

In the input, you must specify only the parameter that you want to update, along with the LUN key. In this case it is the key of the Performance Service Level.

Sample cURL

```
curl -X PATCH "https://<hostname>/api/storage-provider/luns/7d5a59b3-953a-11e8-8857-00a098dcc959" -H "accept: application/json" -H "Content-Type: application/json" -H "Authorization: Basic <Base64EncodedCredentials>" -d "{ \"performance_service_level\": { \"key\": \"1251e51b-069f-11ea-980d-fa163e82bbf2\" } }"
```

The JSON output displays a Job object key that you can use to verify the LUN that you updated.

5. View the details of the LUN by running the following API with the LUN key that you obtained.

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

View the details of the LUN in the output. You can see that the Performance Service Level is assigned to this LUN.

Sample JSON output

```

"assigned_performance_service_level": {
  "key": "1251e51b-069f-11ea-980d-fa163e82bbf2",
  "name": "Value",
  "peak_iops": 75,
  "expected_iops": 75,
  "_links": {
    "self": {
      "href": "/api/storage-provider/performance-service-
levels/1251e51b-069f-11ea-980d-fa163e82bbf2"
    }
  }
}

```

Modifying an NFS file share to support CIFS

You can modify an NFS file share to support CIFS protocol. During file share creation, it is possible to specify both access control list (ACL) parameters and export policy rules for the same file share. However, if you want to enable CIFS on the same volume where you created an NFS file share, you can update the ACL parameters on that file share to support CIFS.

What you'll need

1. An NFS file share must have been created with only the export policy details. For information, see *Managing file shares* and *Modifying storage workloads*.
2. You must have the file share key to run this operation. For information about viewing file share details and retrieving the file share key by using the Job ID, see *Provisioning CIFS and NFS file shares*.

This is applicable for an NFS file share that you created by adding only export policy rules and not ACL parameters. You modify the NFS file share to include the ACL parameters.

Steps

1. On the NFS file share, perform a `PATCH` operation with the ACL details for allowing CIFS access.

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

Sample cURL

Based on the access privileges you assign to the user group, as displayed in the following sample, an ACL is created and assigned to the file share.


```

{
  "access_control": {
    "acl": [
      {
        "permission": "read",
        "user_or_group": "everyone"
      }
    ],
    "active_directory_mapping": {
      "key": "3b648c1b-d965-03b7-20da-61b791a6263c"
    }
  }
}

```

Sample JSON output

The operation returns the Job ID of the Job that runs the update.

2. Verify whether the parameters have been added correctly by querying the file share details for the same file share.

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

Sample JSON output

```

"access_control": {
  "acl": [
    {
      "user_or_group": "everyone",
      "permission": "read"
    }
  ],
  "export_policy": {
    "id": 1460288880641,
    "key": "7d5a59b3-953a-11e8-8857-00a098dcc959:type=export_policy,uuid=1460288880641",
    "name": "default",
    "rules": [
      {
        "anonymous_user": "65534",
        "clients": [
          {
            "match": "0.0.0.0/0"
          }
        ]
      }
    ],
  },
}

```

```

        "index": 1,
        "protocols": [
            "nfs3",
            "nfs4"
        ],
        "ro_rule": [
            "sys"
        ],
        "rw_rule": [
            "sys"
        ],
        "superuser": [
            "none"
        ]
    },
    {
        "anonymous_user": "65534",
        "clients": [
            {
                "match": "0.0.0.0/0"
            }
        ],
        "index": 2,
        "protocols": [
            "cifs"
        ],
        "ro_rule": [
            "ntlm"
        ],
        "rw_rule": [
            "ntlm"
        ],
        "superuser": [
            "none"
        ]
    }
],
"_links": {
    "self": {
        "href": "/api/datacenter/protocols/nfs/export-
policies/7d5a59b3-953a-11e8-8857-
00a098dcc959:type=export_policy,uuid=1460288880641"
    }
}
},

```

```
  "_links": {
    "self": {
      "href": "/api/storage-provider/file-shares/7d5a59b3-953a-11e8-8857-00a098dcc959:type=volume,uuid=e581c23a-1037-11ea-ac5a-00a098dcc6b6"
    }
  }
}
```

You can see the ACL assigned along with the export policy to the same file share.

Copyright information

Copyright © 2024 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—without 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.