



API Developer's Guide



Unified Manager for DM Series 9.7

Second edition (February 2023)

© Copyright Lenovo 2020, 2023.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant to a General Services Administration (GSA) contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925

Contents

Chapter 1. Getting started with Unified Manager for DM Series 1

Audience for this guide	1
Unified Manager for DM Series API access and categories	1
REST services offered in Unified Manager for DM Series	2
API version in Unified Manager for DM Series	3
Storage resources in ONTAP	3

Chapter 2. REST API access and authentication in Unified Manager for DM Series 5

HTTP status codes used in Unified Manager for DM Series	8
Recommendations for using the APIs for Unified Manager for DM Series	9
Logs for troubleshooting	9
Hello API server	10

Chapter 3. Unified Manager REST APIs 13

APIs for managing data centers	13
Gateway APIs	16
Administration APIs	18
Security APIs	19

Job objects API and asynchronous processes	20
Viewing Jobs	21
APIs for managing workloads	21
Viewing storage workloads	21
Managing access endpoints	22
Managing Active Directory mapping	23
Managing file shares	23
Managing LUNs	24
Managing Performance Service Levels	25
Managing Storage Efficiency Policies	26

Chapter 4. Common workflows for storage management 29

Determining space issues in aggregates	29
Workflows for workload management	31
Verifying SVMs on clusters.	31
Provisioning CIFS and NFS file shares	32
Provisioning LUNs	39
Modifying storage workloads.	45

Appendix A. Contacting Support 51

Appendix B. Notices. 53

Trademarks	54
----------------------	----

Chapter 1. Getting started with Unified Manager for DM Series

Unified Manager for DM Series (formerly ThinkSystem Intelligent Monitoring 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.

The Unified Manager *API Developer's Guide* provides you with information about APIs and sample codes. The information provided in the guide enables you to create RESTful clients of Lenovo Manageability software solutions for managing Lenovo systems. The APIs are based on the Representational State Transfer (REST) architectural style.

Unified Manager provides API offerings for storage management in your Lenovo environment. All the four REST operations Create, Read, Update, and Delete (also known as CRUD) are supported.

Audience for this guide

This guide is intended for developers creating applications that interface with the Unified Manager for DM Series software through REST APIs.

The guide is also intended for storage administrators and architects who want to gain a basic understanding of how the REST APIs that are provided in Unified Manager can be used to build client applications to manage and monitor Lenovo storage systems.

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

Note: 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.

Unified Manager for DM Series API access and categories

The Unified Manager for DM Series 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
```

Note: 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/Documentation button → Help/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.

data center	This category contains the APIs that help you to view and manage your data center storage objects. 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 API. Jobs are operations that are scheduled for asynchronous execution related to managing of storage objects or workloads on Unified Manager.
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.

REST services offered in Unified Manager for DM Series

You should be aware of the REST services and operations offered, before you start using the Unified Manager for DM Series 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 /management-server/admin/backup-file-info retrieves backup file details. The maximum number of records that is returned by the GET operation is 1000.

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

- Example for PATCH: PATCH /storage-provider/svms/{key} modifies the properties of an SVM, using its unique key.
- Example for POST: POST /storage-provider/svms creates a custom Storage Virtual Machine (SVM).
- 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.

Table 1.

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.

API version in Unified Manager for DM Series

The REST API URIs in Unified Manager for DM Series 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.

Note: 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 Unified Manager for DM Series, 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, 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:



Chapter 2. REST API access and authentication in Unified Manager for DM Series

The Unified Manager for DM Series REST API is accessible by using any web browser or programming platform that can issue HTTP requests. Unified Manager supports basic HTTP authentication mechanism. Before you call the Unified Manager REST API, you must authenticate a user.

REST access

You can use any web browser or programming platform that can issue HTTP requests to access the Unified Manager REST API. For example, after logging in to Unified Manager, you can type the URL in any browser to retrieve the attributes of all of the management stations, such as the management station name, key, and IP address.

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

Response

```
{
  "records": [
    {
      "key": "65821433-683b-11ea-982c-00a098ddea9d:type=cluster,uuid=65821433-683b-11ea-982c-00a098ddea9d",
      "name": "DM5000_3",
      "uuid": "65821433-683b-11ea-982c-00a098ddea9d",
      "contact": null,
      "location": null,
      "version": {
        "full": "<ONTAP_release>",
        "generation": 9,
        "major": 7,
        "minor": 0
      },
      "isSanOptimized": false,
      "management_ip": "<mgmt._IP>",
      "nodes": [
        {
          "key": "65821433-683b-11ea-982c-00a098ddea9d:type=cluster_node,uuid=50f33ca2-683a-11ea-982c-00a098ddea9d",
          "uuid": "50f33ca2-683a-11ea-982c-00a098ddea9d",
          "name": "DM5000_3-01",
          "_links": {
            "self": {
              "href": "/api/v2/datacenter/cluster/nodes/65821433-683b-11ea-982c-00a098ddea9d:type=cluster_node,uuid=50f33ca2-683a-11ea-982c-00a098ddea9d"
            }
          },
          "location": null,
          "version": {
            "full": "<ONTAP_release>",
            "generation": 9,
            "major": 7,
            "minor": 0
          },
          "model": "DM5000H",
          "uptime": 9170396,
          "serial_number": "<node_SN>"
        },
        {
          "key": "65821433-683b-11ea-982c-00a098ddea9d:type=cluster_node,uuid=cc35d3b1-6839-11ea-9786-00a098ddea9d",
          "uuid": "cc35d3b1-6839-11ea-9786-00a098ddea9d",
          "name": "DM5000_3-02",
          "_links": {
            "self": {
              "href": "/api/v2/datacenter/cluster/nodes/65821433-683b-11ea-982c-00a098ddea9d:type=cluster_node,uuid=cc35d3b1-6839-11ea-9786-00a098ddea9d"
            }
          },
          "location": null,
          "version": {
            "full": "ONTAP_release",
            "generation": 9,
            "major": 7,
            "minor": 0
          },
          "model": "DM5000H",
          "uptime": 8419138,
          "serial_number": "<node_SN>"
        }
      ],
      "_links": {
        "self": {
          "href": "/api/v2/datacenter/cluster/clusters/65821433-683b-11ea-982c-
```

```

00a098ddea9d:type=cluster,uuid=65821433-683b-11ea-982c-00a098ddea9d"
    }
  },
  {
    "key": "f58e263f-8e40-11ea-a0f2-00a098dbe2fe:type=cluster,uuid=f58e263f-8e40-11ea-a0f2-00a098dbe2fe",
    "name": "DM7000F_4",
    "uuid": "f58e263f-8e40-11ea-a0f2-00a098dbe2fe",
    "contact": null,
    "location": null,
    "version": {
      "full": "Data ONTAP release (Lenovo)",
      "generation": 9,
      "major": 7,
      "minor": 0
    },
    "isSanOptimized": false,
    "management_ip": "<mgmt._IP>",
    "nodes": [
      {
        "key": "f58e263f-8e40-11ea-a0f2-00a098dbe2fe:type=cluster_node,uuid=125d784b-8e39-11ea-ae2e-00a098dbe286",
        "uuid": "125d784b-8e39-11ea-ae2e-00a098dbe286",
        "name": "DM7000F_4_2",
        "_links": {
          "self": {
            "href": "/api/v2/datacenter/cluster/nodes/f58e263f-8e40-11ea-a0f2-00a098dbe2fe:type=cluster_node,uuid=125d784b-8e39-11ea-ae2e-00a098dbe286"
          }
        },
        "location": null,
        "version": {
          "full": "Data ONTAP Release (Lenovo) ",
          "generation": 9,
          "major": 7,
          "minor": 0
        },
        "model": "DM7000F",
        "uptime": 1656823,
        "serial_number": "<node_SN>"
      },
      {
        "key": "f58e263f-8e40-11ea-a0f2-00a098dbe2fe:type=cluster_node,uuid=ec84a586-8e37-11ea-a0f2-00a098dbe2fe",
        "uuid": "ec84a586-8e37-11ea-a0f2-00a098dbe2fe",
        "name": "DM7000F_4_1",
        "_links": {
          "self": {
            "href": "/api/v2/datacenter/cluster/nodes/f58e263f-8e40-11ea-a0f2-00a098dbe2fe:type=cluster_node,uuid=ec84a586-8e37-11ea-a0f2-00a098dbe2fe"
          }
        },
        "location": null,
        "version": {
          "full": "Data ONTAP Release (Lenovo)",
          "generation": 9,
          "major": 7,
          "minor": 0
        },
        "model": "DM7000F",
        "uptime": 1649181,
        "serial_number": "<node_SN>"
      }
    ],
    "_links": {
      "self": {
        "href": "/api/v2/datacenter/cluster/clusters/f58e263f-8e40-11ea-a0f2-

```

```

00a098dbe2fe:type=cluster,uuid=f58e263f-8e40-11ea-a0f2-00a098dbe2fe"
    }
  }
},
"num_records": 2,
"total_records": 2,
"_links": {
  "self": {
    "href": "/api/v2/datacenter/cluster/clusters"
  }
}
}
}

```

- *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 POST, PATCH, and DELETE HTTP requests from a web browser, you have to use 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 Unified Manager for DM Series

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

HTTP status code	Status code title	Description
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.
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 GET or GET ALL 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 Unified Manager for DM Series

When using the APIs in Unified Manager for DM Series, 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 Unified Manager for DM Series” 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.
- Avoid querying objects by using a combination of the filter operators wild card (*) and pipe (|). It might retrieve an incorrect number of objects.
- 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.

Table 2.

Log location	Use
<code>/var/log/ocie/access_log.log</code>	Contains all API call details, such as the user name of the user invoking the API, start time, execution time, status, and URL. 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.
<code>/var/log/ocum/ocumserver.log</code>	Contains all API execution logs. You can use this log file to troubleshoot and debug the API calls.

Table 2. (continued)

Log location	Use
/var/log/ocie/server.log	Contains all Wildfly server deployments and start/stop service related logs. 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.
/var/log/ocie/au.log	Contains acquisition unit related logs. 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 Unified Manager for DM Series REST APIs.

Hello API server

The *Hello API server* is a sample program that demonstrates how to invoke a REST API in Unified Manager for DM Series 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 Unified Manager for DM Series. You should be aware of the following considerations about the sample code:

- Uses an HTTPS connection to Unified Manager for DM Series to invoke the specified REST URI
- Ignores the certificate provided by Unified Manager for DM Series
- 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 {
URLConnection 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 HttpsURLConnection 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() {
HttpsURLConnection 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());
HttpsURLConnection.setDefaultSSLSocketFactory(sslContext.getSocketFactory());
URL url = new URL(server_url);
conn = (HttpsURLConnection) 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;
}
}

```

Chapter 3. Unified Manager REST APIs

The REST APIs for Unified Manager for DM Series 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.

APIs for managing data centers

The REST APIs under the datacenter category provide information about the clusters, nodes, aggregates, volumes, LUNs, fileshares, namespaces, and other elements in your data center. These APIs are available for querying, adding, deleting, or modifying the configuration in your data center.

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

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

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

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

HTTP verb	Path	Description
GET	/datacenter/cluster/clusters /datacenter/cluster/clusters/{key}	You can use this method to view the details of the ONTAP clusters across the data center.
GET	/datacenter/cluster/nodes /datacenter/cluster/nodes/{key}	You can use this method to view the details of the nodes in the data center.
GET	/datacenter/protocols/cifs/shares /datacenter/protocols/cifs/shares/ {key}	You can use this method to view the details of the CIFS shares in the data center.

HTTP verb	Path	Description
GET	/datacenter/protocols/nfs/export-policies /datacenter/protocols/nfs/export-policies/{key}	You can use this method to view the details of the export policies for the supported NFS services. You can query the export policies for a cluster or storage VM and reuse the export policy key for provisioning NFS file shares. For more information about the assigning and reusing export policies on workloads, see “Provisioning CIFS and NFS file shares”.
GET	/datacenter/storage/aggregates /datacenter/storage/aggregates/{key}	You can use this method to view the collection of aggregates in the entire data center for provisioning workloads on them or monitoring.
GET	/datacenter/storage/luns /datacenter/storage/luns/{key}	You can use this method to view the collection of LUNs in the entire data center.
GET	/datacenter/storage/qos/policies /datacenter/storage/qos/policies/{key}	You can use this method to view the details of all the QoS policies in the data center and apply the policies as required.
GET	/datacenter/storage/qtrees /datacenter/storage/qtrees/{key}	You can use this method to view the qtree details across the data center.
GET	/datacenter/storage/volumes /datacenter/storage/volumes/{key}	You can use this method to view the collection of volumes in the data center.

HTTP verb	Path	Description
GET POST DELETE PATCH	/datacenter/protocols/san/igroups /datacenter/protocols/san/ igroups/{key}	<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. You can use these methods to query, create, delete, and modify igroups respectively. 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.
GET POST DELETE PATCH	/datacenter/svm/svms /datacenter/svm/svms/{key}	<p>You can use these methods to view, create, delete, and modify Storage Virtual Machines (storage VMs). Points to note:</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. <p>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, SMB/CIFS, NFS, FCP, and iSCSI. The provisioning workflows might fail if the storage VM does not support the required services. It is recommended that the services for the respective types of workloads are also enabled on the storage VM.</p> <ul style="list-style-type: none"> • DELETE: You need to provide the storage VM key to delete a particular storage VM. <p>Note: 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/SMB server has been configured, this API also deletes the CIFS/SMB server, along with the local Active Directory configuration. However, the CIFS/SMB server name continues to be in the Active Directory configuration that you must delete manually from the Active Directory server.</p> <ul style="list-style-type: none"> • PATCH: You need to provide the storage VM key to modify a particular storage VM. You also

HTTP verb	Path	Description
		need to enter the properties that you want to update, along with their values.

Gateway APIs

The gateway APIs provide you with the advantage of using the Unified Manager for DM Series 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. The other features are available only 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.

Note: Users with Application Administrator role can run all ONTAP REST operations by using these APIs. Users of all roles can run the `GET` operation.

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	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).
GET POST DELETE PATCH OPTIONS (not available on Swagger) HEAD (not available on Swagger)	/gateways/{uuid}/{path}	<p>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. Tunneling or proxy functionality cannot be disabled.</p> <p>The OPTIONS method returns all the operations supported by an ONTAP REST API. For example, if an ONTAP API supports only the GET operation, running the OPTIONS method by using this gateway API returns <i>GET</i> as the response. This method is not supported on Swagger, but can be performed on other API tools.</p> <p>The OPTIONS 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.</p> <p>Note: The value for {uuid} must be replaced with the cluster UUID on which the REST operation is to be performed. The {path} must be replaced by the ONTAP REST URL. You must remove <i>/api/</i> from the path.</p>

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.

Note: While adding the path, ensure that you have removed the “/” symbol at the beginning of the URL. For the API `/storage/volumes`, add `storage/volumes`.

The appended URL is:

```
/gateways/{uuid}/storage/volumes
```

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

```
GET https://<hostname>/api/gateways/<cluster_UUID>/storage/volumes
```

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

Sample cURL command

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

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

Setting API scope

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

Sample cURL command

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

For more information about using ONTAP REST APIs, see the ONTAP REST API Developers guide.

[ONTAP REST API Developers guide](#)

Administration APIs

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 Unified Manager for DM Series datasources.

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

HTTP verb	Path	Description
GET PATCH	/admin/backup-settings /admin/backup-settings	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: <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 The time of the backup is in server time zone. 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.
GET	/admin/backup-file-info	A backup dump file is generated every time the backup schedule is modified for Unified Manager. You can use this method to verify whether the backup file is generated according to the modified backup settings, and whether the information on the file matches the modified settings.
GET	/admin/datasource-certificate	You can use this method to view the datasource (cluster) certificate from the trust store. Validating the certificate is required before adding an ONTAP cluster as a Unified Manager datasource.
GET POST PATCH DELETE	/admin/datasources/clusters /admin/datasources/clusters/{key}	You can use the <code>GET</code> method to retrieve the details of the datasources (ONTAP clusters) managed by Unified Manager. 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. For modifying and deleting a cluster managed as a datasource by Unified Manager, use the ONTAP cluster key.

Security APIs

You can use the APIs in the security category to control user access to selected cluster objects in Unified Manager for DM Series. 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.

Note: 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 POST	/security/users /security/users	You can use these methods to get the details of users or add new user to Unified Manager. 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.
GET PATCH DELETE	/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.

Job objects API and asynchronous processes

The management-server category consists of the jobs API that provides information about the Jobs performed while running Unified Manager for DM Series 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.

Viewing Jobs

In Unified Manager for DM Series, operations like 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.

You can query all of the Job objects. You can also use the Job key and Job object details to run the next set of operations on the resources.

View jobs

You can use this method to obtain a list of all of the Jobs. The response body consists of the Job details of all of the Jobs. You can also view the details of a particular Job object by specifying the Job key in the URI. The response body consists of the details that are identified by the Job key. In the response, the latest Job object is returned to the top. If you want to query a particular Job object, enter the Job ID of the job.

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

APIs for 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 Unified Manager for DM Series 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 Unified Manager for DM Series 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 Storage Manager or ONTAP CLI to create an access endpoint on each node. Both IPv4 and IPv6 addresses are supported for access endpoints creation.

Note: 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 APIs listed here to view, add, and delete CIFS shares and NFS file shares. Using these APIs, you can also assign and modify the Performance Service Level and Storage Efficiency Policies for the file shares.

Before assigning storage, the virtual machine must have been created and provisioned with the supported protocols. Similarly, for assigning Storage Efficiency Policies to the file shares, the Storage Efficiency Policies must be created before creating the file shares.

View file shares

You can use the following method to view all of the storage workloads available in your Unified Manager environment. When you have added an ONTAP cluster as a datasource on Unified Manager for DM Series, the storage workloads for those clusters are automatically added to your Unified Manager instance. This API retrieves the all 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 creating FlexGroups.

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

Note: 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.

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: 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 APIs listed here to view, add, and delete LUNs on the Storage Virtual Machines (storage VMs). You can also assign and modify the Performance Service Level and Storage Efficiency Policies for the LUNs.

Before assigning storage, you must ensure that the SVM has been created and provisioned with the supported protocols. Similarly, for assigning Performance Service Levels to the LUNs, the Performance Service Levels must be created before creating the LUN.

View LUNs

You can use the following method to view all the LUNs in your Unified Manager environment. When you have added an ONTAP cluster as a datasource on Unified Manager for DM Series, the storage workloads for those clusters are automatically added to your Unified Manager instance. This API retrieves the all 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

Note: 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.

Note: 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”.

Note: 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 Unified Manager for DM Series.

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 Performance Service Level 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.

Notes: 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}

Chapter 4. Common workflows for storage management

The common workflows provide client application developers with examples of how Unified Manager for DM Series 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.

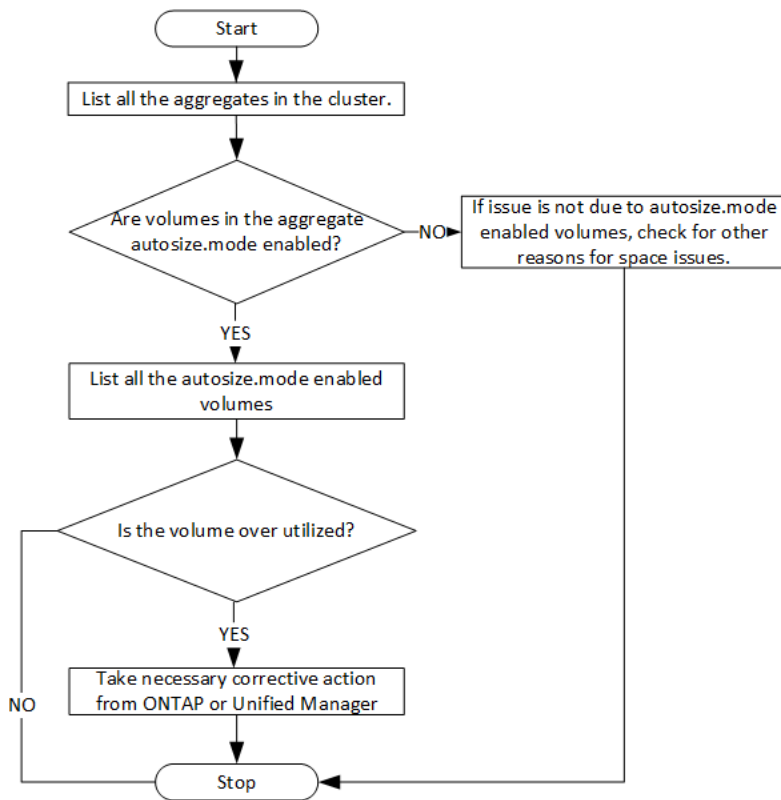
Determining space issues in aggregates

You can use the data center APIs in Unified Manager for DM Series 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.

Workflows for workload management

Using Unified Manager for DM Series, 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.

Note: 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

Note: 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.

Note: 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.

Note: 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-aelc-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-aelc-123478563412,type=object type, uuid=1cd8a442-86d1-11e0-aelc-9876567890123\" }, \"dns\": { \"domains\": [ \"example.com\", \"example2.example3.com\" ], \"servers\": [ \"10.224.65.20\", \"2001:db08:a0b:12f0::1\" ] }, \"fcp\": { \"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.

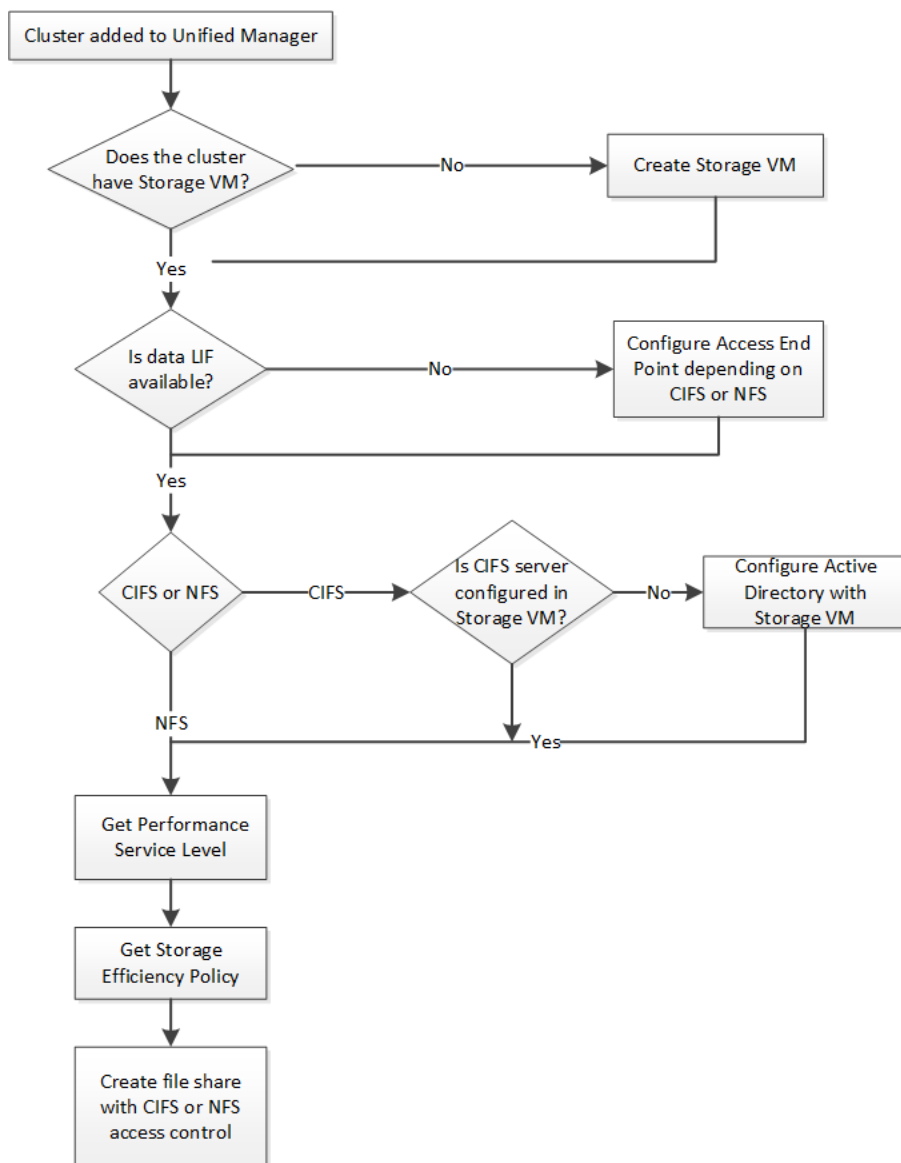
- Verify the SVM creation by using the job object key for query. If the SVM is created successfully, the SVM key is returned in the response.

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 Unified Manager for DM Series. 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.



Notes: 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>"
```

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

Note: 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.

- Verify the access endpoint:

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

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

- 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

- 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

Note: 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.

Note: 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:

Note: 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.

9. 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"
    }
  }
}

```

10. 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": "DM7000F-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": [

```

```

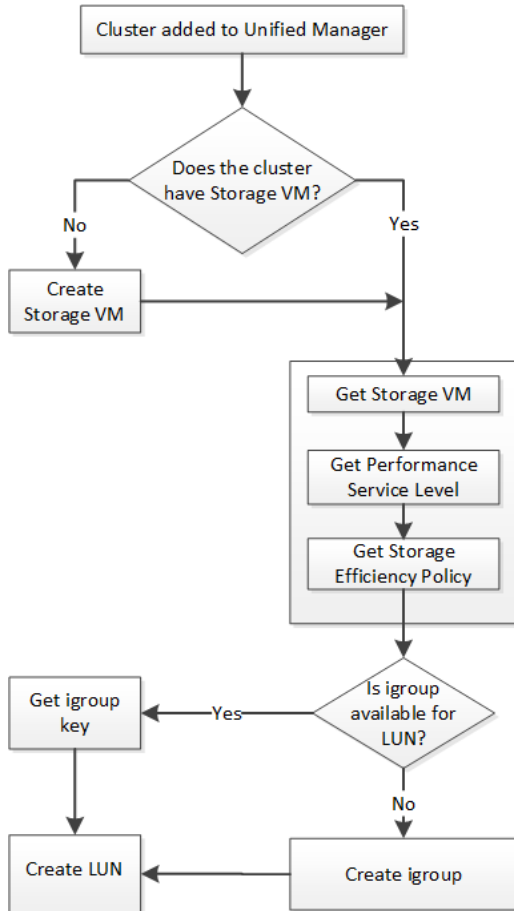
    {
      "user_or_group": "everyone",
      "permission": "read"
    }
  ],
  "export_policy": {
    "id": 1460288880641,
    "key": "7d5a59b3-953a-11e8-8857-00a098dcc959:type=export_policy,uuid=1460288880641",
    "name": "default",
    "rules": [
      {
        "anonymous_user": "65534",
        "clients":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

```

Provisioning LUNs

You can provision LUNs on your Storage Virtual Machines (SVMs) by using the provisioning APIs provided as a part of Unified Manager for DM Series. This provisioning workflow details the steps for retrieving the keys of the SVMs, Performance Service Levels, and Storage Efficiency Policies before creating the LUN.

The following diagram illustrates the steps in a LUN provisioning workflow.



Note: 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

Note: 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 igroup 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.

Note: 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.

```
curl -X POST "https://<hostname>/api/storage-provider/luns" -H "accept: application/json" -H "Content-Type: application/json" -H "Authorization: Basic <Base64EncodedCredentials>" -d '{"name": "MigrationLunWithVol", "os_type": "windows", "performance_service_level": { "key": "7873dc0d-0ee5-11ea-82d7-fa163ea0eb69" }, "space": { "size": 1024000000 }, "svm": { "key": "333fbcfa-0ace-11ea-9d6d-00a09897cc15:type=vserver,uuid=4d462ec8-0f56-11ea-9d6d-00a09897cc15" } }'
```

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.

```

    {
      "name": "lunKey",
      "value": "key": "f963839f-0f95-11ea-9963-00a098884af5:type=lun,uuid=71f3187e-bf19-4f34-ba34-b1736209b45a"
    }
  ],
  "_links": {
    "self": {
      "href": "/api/management-server/jobs/fa7c856d29e2b80f%3A-8d3325d%3A16e9eb5ed6d%3A-548b"
    }
  }
}

```

8. 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.

```

{
  "key": "f963839f-0f95-11ea-9963-00a098884af5:type=lun,uuid=71f3187e-bf19-4f34-ba34-b1736209b45a",
  "name": "/vol/NSLM_VOL_LUN_1574753881051/LunForTesting1",
  "uuid": "71f3187e-bf19-4f34-ba34-b1736209b45a",
  "cluster": {
    "uuid": "f963839f-0f95-11ea-9963-00a098884af5",
    "key": "f963839f-0f95-11ea-9963-00a098884af5:type=cluster,uuid=f963839f-0f95-11ea-9963-00a098884af5",
    "name": "sti2552-4451574693410",
    "_links": {
      "self": {
        "href": "/api/datacenter/cluster/clusters/f963839f-0f95-11ea-9963-00a098884af5:type=cluster,uuid=f963839f-0f95-11ea-9963-00a098884af5"
      }
    }
  },
  "svm": {
    "uuid": "7754a99c-101f-11ea-9963-00a098884af5",
    "key": "f963839f-0f95-11ea-9963-00a098884af5:type=vserver,uuid=7754a99c-101f-11ea-9963-00a098884af5",
    "name": "Testingsvml",
    "_links": {
      "self": {
        "href": "/api/datacenter/svm/svms/f963839f-0f95-11ea-9963-00a098884af5:type=vserver,uuid=7754a99c-101f-11ea-9963-00a098884af5"
      }
    }
  },
  "volume": {
    "uuid": "961778bb-2be9-4b4a-b8da-57c7026e52ad",
    "key": "f963839f-0f95-11ea-9963-00a098884af5:type=volume,uuid=961778bb-2be9-4b4a-b8da-57c7026e52ad",
    "name": "NSLM_VOL_LUN_1574753881051",
    "_links": {
      "self": {
        "href": "/api/datacenter/storage/volumes/f963839f-0f95-11ea-9963-00a098884af5:type=volume,uuid=961778bb-2be9-4b4a-b8da-57c7026e52ad"
      }
    }
  },
  "assigned_performance_service_level": {
    "key": "861f6e4d-0c35-11ea-9d73-fa163e706bc4",
    "name": "Value",
    "peak_iops": 75,
    "expected_iops": 75,
  }
}

```

```

    "_links": {
      "self": {
        "href": "/api/storage-provider/performance-service-levels/861f6e4d-0c35-11ea-9d73-fa163e706bc4"
      }
    }
  },
  "recommended_performance_service_level": {
    "key": null,
    "name": "Idle",
    "peak_iops": null,
    "expected_iops": null,
    "_links": {}
  },
  "assigned_storage_efficiency_policy": {
    "key": null,
    "name": "Unassigned",
    "_links": {}
  },
  "space": {
    "size": 1024458752
  },
  "os_type": "linux",
  "_links": {
    "self": {
      "href": "/api/storage-provider/luns/f963839f-0f95-11ea-9963-00a098884af5%3Atype%3Dlun%2Cuuid%3D71f3187e-bf19-4f34-ba34-b1736209b45a"
    }
  }
}

```

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.

Note: 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.

```

curl -X POST "https://<hostname>/api/storage-provider/access-endpoints" -H "accept: application/json" -H "Content-Type: application/json" -H "Authorization: Basic <Base64EncodedCredentials>" -d "{ \"data_protocols\": [ \"iscsi\" ], \"ip\": { \"address\": \"10.162.83.126\", \"ha_address\": \"10.142.83.126\", \"netmask\": \"255.255.0.0\" }, \"lun\": { \"key\": \"e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=lun,uuid=b8e0clae-0997-47c5-97d2-1677d3ec08ff\" }, \"name\": \"aep_example\" }"

```

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**.

```

"job_results": [
  {
    "name": "accessEndpointKey",
    "value": "e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=network_lif,lif_uuid=9c964258-14ef-11ea-95e2-00a098e32c28"
  },
  {
    "name": "accessEndpointHAKey",
    "value": "e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=network_lif,lif_uuid=9d347006-14ef-11ea-8760-00a098e3215f"
  }
],
"_links": {
  "self": {
    "href": "/api/management-server/jobs/71377eeee0b25633%3A-30a2dbfe%3A16ec620945d%3A-7f5a"
  }
}
}

```

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

```

curl -X PATCH "https://<hostname>/api/storage-provider/luns/e4f33f90-f75f-11e8-9ed9-00a098e3215f%3Atype%3Dlun%2Cuuid%3Db8e0c1ae-0997-47c5-97d2-1677d3ec08ff"
-H "accept: application/json" -H "Content-Type: application/json" -H "Authorization: Basic
<Base64EncodedCredentials>" -d
"{ \"lun_maps\": [ { \"igroup\":
{ \"key\": \"e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=igroup,uid=d19ec2fa-fec7-11e8-b23d-00a098e32c28\" },
\"logical_unit_number\": 3 } ]}"

```

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.

```
{
  "key": "e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=lun, uuid=b8e0c1ae-0997-47c5-97d2-1677d3ec08ff",
  "name": "/vol/NSLM_VOL_LUN_1575282642267/example_lun",
  "uuid": "b8e0c1ae-0997-47c5-97d2-1677d3ec08ff",
  "cluster": {
    "uuid": "e4f33f90-f75f-11e8-9ed9-00a098e3215f",
    "key": "e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=cluster, uuid=e4f33f90-f75f-11e8-9ed9-00a098e3215f",
    "name": "umeng-dm5000f-01-02",
    "_links": {
      "self": {
        "href": "/api/datacenter/cluster/clusters/e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=cluster,
        uuid=e4f33f90-f75f-11e8-9ed9-00a098e3215f"
      }
    }
  },
  "svm": {
    "uuid": "97f47088-fa8e-11e8-9ed9-00a098e3215f",
    "key": "e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=vserver, uuid=97f47088-fa8e-11e8-9ed9-00a098e3215f",
    "name": "NSLM12_SVM_ritu",
    "_links": {
      "self": {
        "href": "/api/datacenter/svm/svms/e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=vserver,
        uuid=97f47088-fa8e-11e8-9ed9-00a098e3215f"
      }
    }
  },
  "volume": {
    "uuid": "a1e09503-a478-43a0-8117-d25491840263",
    "key": "e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=volume, uuid=a1e09503-a478-43a0-8117-d25491840263",
    "name": "NSLM_VOL_LUN_1575282642267",
    "_links": {
      "self": {
        "href": "/api/datacenter/storage/volumes/e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=volume,
        uuid=a1e09503-a478-43a0-8117-d25491840263"
      }
    }
  },
  "lun_maps": [
    {
      "igroup": {
        "uuid": "d19ec2fa-fec7-11e8-b23d-00a098e32c28",
        "key": "e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=igroup, uuid=d19ec2fa-fec7-11e8-b23d-00a098e32c28",
        "name": "lun55_igroup",
        "_links": {
          "self": {
            "href": "/api/datacenter/protocols/san/igroups/e4f33f90-f75f-11e8-9ed9-00a098e3215f:type=igroup,
            uuid=d19ec2fa-fec7-11e8-b23d-00a098e32c28"
          }
        }
      }
    }
  ],
  "logical_unit_number": 3
},
"assigned_performance_service_level": {
  "key": "cf2aacda-10df-11ea-bbe6-fa163e599489",
  "name": "Value",
  "peak_iops": 75,
  "expected_iops": 75,
  "_links": {
    "self": {
      "href": "/api/storage-provider/performance-service-levels/cf2aacda-10df-11ea-bbe6-fa163e599489"
    }
  }
},
"recommended_performance_service_level": {
  "key": null,

```



```

    "name": "Idle",
    "peak_iops": null,
    "expected_iops": null,
    "_links": {}
  },
  "assigned_storage_efficiency_policy": {
    "key": null,
    "name": "Unassigned",
    "_links": {}
  },
  "space": {
    "size": 1073741824
  },
  "os_type": "linux",
  "_links": {
    "self": {
      "href": "/api/storage-provider/luns/e4f33f90-f75f-11e8-9ed9-00a098e3215f%3Atype%3Dlun%2Cuuid%3Db8e0c1ae-0997-47c5-97d2-1677d3ec08ff"
    }
  }
}

```

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.

Note: 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.

Note: 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": {}
},
```

- 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

Note: 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.

- 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.

- 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

Note: 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

Note: 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.

Before you begin

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*.

About this task

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.

Procedure

Step 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.

Step 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.

Modifying workloads to upgrade QoS to AQoS

Unified Manager for DM Series provides support for both traditional QoS and Adaptive QoS (AQoS) that are associated with the storage workloads managed by Unified Manager.

Appendix A. Contacting Support

You can contact Support to obtain help for your issue.

You can receive hardware service through a Lenovo Authorized Service Provider. To locate a service provider authorized by Lenovo to provide warranty service, go to <https://datacentersupport.lenovo.com/serviceprovider> and use filter searching for different countries. For Lenovo support telephone numbers, see <https://datacentersupport.lenovo.com/supportphonenumberlist> for your region support details.

Appendix B. Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area.

Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document is not an offer and does not provide a license under any patents or patent applications. You can send inquiries in writing to the following:

*Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing*

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Trademarks

LENOVO, LENOVO logo, and THINKSYSTEM are trademarks of Lenovo. All other trademarks are the property of their respective owners. © 2023 Lenovo.

Lenovo