

Dell EMC OpenManage Enterprise-Modular Edition Version 1.10.00 for PowerEdge MX7000 Chassis

User's Guide

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

1 Overview	8
Key features	8
New in this release	8
Supported platforms	9
Supported web browsers	9
Other documents you may need	9
Accessing documents from Dell support site	10
Positioning OME-Modular with other Dell EMC applications	10
2 Updating the management module firmware	11
Updating the firmware using catalog-based compliance method	11
3 Logging in to OME-Modular	12
Logging in to OME-Modular as local, Active Directory, or LDAP user	12
Logging in to OME-Modular as Active Directory or LDAP user	13
OME-Modular home page	13
Multi-chassis management dashboard	14
Viewing device health	15
Setting up chassis	15
Initial configuration	15
Configuring chassis settings	16
Configure chassis power	16
Configure chassis management network	17
Configure chassis network services	19
Configure local access	19
Configure chassis location	21
Configure Quick Deploy settings	22
Managing chassis	23
Creating chassis filters	23
Viewing chassis overview	23
Wiring chassis	24
Chassis groups	25
Prerequisites for creating a wired group	25
Creating chassis groups	26
MCM dashboard	28
Controlling chassis power	28
Backing up chassis	29
Restoring chassis	29
Exporting chassis profiles	30
Managing chassis failover	30
Troubleshooting in chassis	30
Blinking LEDs	30
Interfaces to access OME-Modular	31
Viewing chassis hardware	32

Viewing chassis alerts.....	32
Viewing chassis hardware logs.....	33
Configuring OME–Modular.....	33
Viewing current configuration.....	33
Configuring users and user settings.....	36
Configuring login security settings.....	40
Configuring alerts.....	41
4 Managing compute sleds.....	43
Viewing compute overview.....	43
Configuring compute settings.....	45
Configuring compute network settings.....	45
Replacing compute sleds.....	45
Viewing compute hardware.....	46
Viewing compute firmware.....	46
Viewing compute hardware logs.....	47
Viewing compute alerts.....	47
5 Managing storage.....	48
Storage overview.....	48
Viewing hardware details.....	49
Assigning drives to a compute sled.....	50
Assigning storage enclosure to a compute sled.....	50
Replacing storage sleds.....	51
Updating enclosure firmware.....	51
Updating the firmware using DUP.....	51
Updating the firmware using catalog-based compliance.....	51
Downgrading storage enclosure firmware.....	51
Managing SAS IOMs.....	51
SAS IOM Overview.....	52
Force active.....	53
Clearing configuration.....	53
Extracting IOM logs.....	53
6 Managing templates.....	54
Viewing template details.....	54
Creating templates.....	54
Importing templates.....	55
Deploying templates.....	55
Deploying templates from Template Details page.....	55
Editing templates.....	55
Editing template networks.....	56
Cloning templates.....	56
Exporting templates.....	56
Deleting templates.....	56
7 Managing identity pools.....	57
Creating identity pools.....	57
Editing identity pools.....	58

Exporting identity pools.....	59
Deleting identity pools.....	59
8 Ethernet IO Modules.....	60
Viewing hardware details.....	61
Configuring IOM settings.....	61
Configuring IOM network settings.....	61
Configuring Linux administrator password.....	62
Configuring SNMP settings.....	62
Configuring advanced settings.....	62
Configuring ports.....	62
9 MX Scalable Fabric architecture.....	64
Recommended physical topology.....	64
Restrictions and guidelines.....	65
Recommended connection order.....	66
10 SmartFabric Services.....	67
Guidelines for operating in SmartFabric mode.....	68
SmartFabric network topologies.....	68
Switch to switch cabling.....	69
Upstream network switch requirements.....	70
NIC teaming restrictions.....	70
CLI commands available in Fabric mode.....	71
Viewing fabric details.....	71
Adding fabric.....	71
Adding uplinks.....	72
Adding network.....	72
Editing uplink.....	73
Viewing topology details.....	73
Editing fabric details.....	73
Deleting uplinks.....	73
Deleting fabric.....	74
VLANs for SmartFabrics and FCoE.....	74
Defining VLANs.....	74
Editing VLANs.....	74
VLAN scaling guidelines.....	75
11 Managing networks.....	76
SmartFabric VLAN management and automated QoS.....	76
Defining networks.....	77
Editing networks.....	77
Exporting network configurations.....	77
Deleting network configurations.....	78
12 Managing Fibre Channel IOMs.....	79
13 Managing firmware.....	80
Creating baselines.....	80

Checking compliance.....	81
Editing baselines.....	81
Managing catalogs.....	81
Viewing catalogs.....	82
Adding catalogs.....	82
Updating firmware.....	83
Rolling back firmware.....	83
Deleting firmware.....	84
14 Monitoring alerts and logs.....	85
Alert log.....	85
Filtering alert logs.....	85
Acknowledging alert logs.....	86
Unacknowledging alert logs.....	86
Ignoring alert logs.....	86
Exporting alert logs.....	86
Deleting alert logs.....	86
Alert policies.....	86
Creating alert policies.....	87
Enabling alert policies.....	87
Editing alert policies.....	88
Disabling alert policies.....	88
Deleting alert policies.....	88
Alert definitions.....	88
Filtering alert definitions.....	88
15 Monitoring audit logs.....	90
Filtering audit logs.....	90
Exporting audit logs.....	90
Monitoring jobs.....	91
Filtering jobs.....	91
Viewing job details.....	92
Running jobs.....	92
Stopping jobs.....	92
Enabling jobs.....	93
Disabling jobs.....	93
Deleting jobs.....	93
16 Use case scenarios.....	94
Assigning backup to the MCM Lead.....	94
Creating chassis group with backup lead.....	94
Monitoring the MCM group.....	95
Scenarios when backup lead can take over as lead chassis.....	96
Disaster recovery of lead chassis.....	96
Retiring lead chassis.....	98
17 Troubleshooting.....	99
Storage.....	99
Firmware update is failing.....	99

Storage assignment is failing.....	99
SAS IOM status is downgraded.....	99
SAS IOM health is downgraded.....	99
Drives on compute sled are not visible.....	100
Storage configuration cannot be applied to SAS IOMs.....	100
Drives in OpenManage are not visible.....	100
iDRAC and OpenManage drive information do not match.....	100
The assignment mode of storage sled is unknown.....	100
Unable to access OME-Modular using Chassis Direct.....	100
Troubleshooting lead chassis failure.....	100
A Recommended slot configurations for IOMs.....	101
Supported slot configurations for IOMs.....	101

Overview

The Dell EMC OpenManage Enterprise Modular (OME-Modular) application runs on the PowerEdge M9002m management module (MM) firmware. OME-Modular facilitates configuration and management of a standalone PowerEdge MX chassis or group of MX chassis using a single Graphical User Interface (GUI). You can use OME-Modular to deploy servers and update firmware. You can also manage the overall health of the chassis and the chassis components such as compute sleds, network devices, input or output modules (IOMs), and storage devices. OME-Modular also facilitates the following activities on the hardware:

- Connectivity of management network.
- Discovery and inventory.
- Monitoring and power control operations and thermal functions.

You can use OME-Modular to manage key workloads on the MX7000 platforms.

- Large and unstructured data and analytics
- Hyper converged and traditional workloads
- Database workloads
- Software defined storage
- HPC and performance workloads

The lead chassis in the Multi Chassis Management (MCM) enables you to perform the following tasks:

- Manage servers across multiple MX chassis.
- Deploy or update servers from lead chassis without launching the member chassis web interface.
- Manage fabric switch engines in fabric mode using the OME-Modular web interface.
- Manage alert log and actions.
- Manage virtual MAC/WWN identity pools.
- Deploy compute sleds easily using server profiles and templates.

OME-Modular offers simple and static roles such as the chassis administrator, compute manager, fabric manager, storage manager, and viewer roles while, OpenManage Enterprise offers static and dynamic groups with role-based access control (RBAC).

Topics:

- [Key features](#)
- [New in this release](#)
- [Supported platforms](#)
- [Supported web browsers](#)
- [Other documents you may need](#)
- [Accessing documents from Dell support site](#)
- [Positioning OME-Modular with other Dell EMC applications](#)

Key features

The key features of OME-Modular are:

- End-to-end life cycle management for servers, storage, and networking.
- Addition of a new chassis to add server, storage, and networking capacity.
- Multiple chassis management using a single interface—web or RESTful interface.
- Management of network IOMs and SmartFabric Services.
- Usage of the automation and security features of iDRAC9.

New in this release

This release of OME-Modular supports:

- Group management through LCD
- Assigning member chassis in MCM as backup chassis and promoting backup chassis as lead chassis

- Configuring hot spare for the chassis
- Accessing chassis using USB
- Customizing strings that are displayed on the chassis LCD
- Configuring network session inactivity timeout
- Extracting logs to a local drive on your system
- Downloading the Management Information Base (MIB) file to a local drive on your system.

Supported platforms

OME - Modular supports the following platforms and components:

Platforms:

- PowerEdge MX7000
- PowerEdge MX740c
- PowerEdge MX840c
- PowerEdge MX5016s
- PowerEdge MX5000s SAS Switch
- PowerEdge MX 25 Gb Ethernet Pass-Through Module
- MX 10GBASE-T Ethernet Pass-Through Module
- Dell EMC MX9116n Fabric Switching Engine
- Dell EMC MX5108n Ethernet Switch
- Dell EMC MX7116n Fabric Expander Module
- Dell EMC MXG610s Fibre Channel Switching Module
- PowerEdge MX9002m Management module

Supported web browsers

OME-Modular is supported on the following web browsers:

- Google Chrome version 63
- Google Chrome version 64
- Mozilla Firefox version 57
- Mozilla Firefox version 58
- Microsoft EDGE
- Microsoft Internet Explorer 11
- Safari version 11

For the OME-Modular web interface to load properly in the web browsers, ensure that the Active X/Java script and font download options are enabled.

 **NOTE: OME-Modular supports TLS 1.2 and later versions.**

Other documents you may need

For more information about managing your system, access the following documents:

Table 1. List of other documents for reference


Name of the document	Brief introduction of the document
<i>OpenManage Enterprise Modular RACADM Command Line Reference Guide</i>	This document contains information about the RACADM subcommands, supported interfaces, and property database groups and object definitions.
<i>OpenManage Enterprise Modular Release Notes</i>	This document provides the latest updates to the system or documentation or advanced technical reference material that is intended for experienced users or technicians.

Table 1. List of other documents for reference(continued)

Name of the document	Brief introduction of the document
OpenManage Enterprise and OpenManage Enterprise – Modular RESTful API Guide	This document provides information about integrating your applications with OpenManage Enterprise Modular, using the RESTful API commands.
<i>Integrated Dell Remote Access Controller (iDRAC) User's Guide</i>	This document provides information about installation, configuration, and maintenance of the iDRAC on managed systems.
<i>OS10 Enterprise Edition User Guide</i>	This document provides information about the features of the OS10 switches and using commands in the IOM CLI to configure the switches.
<i>PowerEdge MX SmartFabric Configuration and Troubleshooting Guide</i>	This document provides information about configuring and troubleshooting SmartFabric Services running on PowerEdge MX systems.
<i>Dell EMC PowerEdge MX7000 Enclosure Installation and Service Manual</i>	This document provides information about installing and replacing components in the PowerEdge MX7000 enclosure.
<i>Dell EMC PowerEdge MX5016s and MX5000s Installation and Service Manual</i>	This document provides information about installing and replacing components in the PowerEdge MX5016s storage sled and PowerEdge MX5000s SAS IOM.

Accessing documents from Dell support site

You can access the required documents in one of the following ways:

- Using the following links:
 - For OpenManage documents — www.dell.com/openmanagemanuals
 - For iDRAC and Lifecycle Controller documents — www.dell.com/idracmanuals
 - For all Enterprise Systems Management documents — www.dell.com/esmmanuals
 - For OpenManage Connections Enterprise Systems Management documents — www.dell.com/esmmanuals
 - For Serviceability Tools documents — <https://www.dell.com/serviceabilitytools>
 - For Client Command Suite Systems Management documents — www.dell.com/omconnectionsclient
 - For SmartFabric OS10 documents—
- From the Dell Support site:
 1. Go to <https://www.dell.com/support>.
 2. Click **Browse all products**.
 3. Click the desired product category, such as Servers, Software, Storage, and so on.
 4. Click the desired product and then click the desired version if applicable.
 **NOTE: For some products, you may need to navigate through the subcategories.**
 5. Click **Manuals & documents**.

Positioning OME-Modular with other Dell EMC applications

OME-Modular works with the following applications to manage, simplify, and streamline operations:

- OME-Modular discovers and inventories MX 7000 chassis in the data center using the OME-Modular REST API commands.
- integrated Dell Remote Access Controller (iDRAC)—OME-Modular manages virtual consoles through iDRAC.
- Repository Manager—OME-Modular uses Repository Manager to create custom repositories in shared networks for creating catalogs. The catalogs are used for firmware updates.
- OME-Modular extracts the OpenManage SupportAssist logs from iDRAC for resolving issues.

Updating the management module firmware

In MCM environment, perform the firmware update for all devices from the lead chassis. Also, select the IOMs and storage sleds as individual devices and not as chassis components, for a successful firmware update.

NOTE: Ensure that you upgrade the OME-Modular firmware to 1.10.00 before upgrading to OS10.

You can update the management module firmware using the following methods:

1. Individual package method—Through OME-Modular web interface or RESTful API
2. Catalog-based compliance method

To update the firmware using the Individual package method:

1. Download the DUP from the www.dell.com/support/drivers.
2. On the OME-Modular web interface, go to **Devices > Chassis** and select the chassis for which you want to update the firmware.
3. Click **Update Firmware**.
The **Select Firmware Source** window is displayed.
4. Select the **Individual package** option and click **Browse** to go to the location where you have downloaded the DUP and click **Next**.
Wait for the comparison report. The supported components are displayed.
5. Select the required components, for example: OME-Modular, and click **Update** to start the firmware update.
You can schedule the update process to start at the time you want.
6. Navigate to the **Monitor > Jobs** page to view the job status.

NOTE: The console is inaccessible during the OME-Modular update process. After the OME-Modular update process, allow 3-5 minutes for the console to reach a steady state.

Topics:

- [Updating the firmware using catalog-based compliance method](#)

Updating the firmware using catalog-based compliance method

To update the firmware using the catalog-based compliance method:

1. Download the DUP from the www.dell.com/support/drivers.
2. Use the Dell Repository Manager (Repository Manager) to create the `catalog.xml` file.
3. Place the `catalog.xml` that you created using Repository Manager, in a shared location.
4. Go to the **Configuration Firmware** page to create the catalog and baseline.
5. In the OME-Modular web interface, go to the **Devices > Chassis** page.
6. Click **Update Firmware** option. The **Select Firmware Source** window is displayed.
7. Select the **Baseline** option and select the required baseline from the drop-down.
8. Select the OME-Modular component from the comparison report.
The supported components are displayed.
9. Select the required components, for example: OME-Modular, and click **Update** to start the firmware update.
10. Navigate to the **Monitor > Jobs** page to view the job status.

NOTE: Use the **Add** option on the **Configuration > Firmware > Catalog Management** option to download the catalog from the <https://www.dell.com/support>.

Logging in to OME-Modular

You can log in to OME-Modular as a local, Active Directory, or generic LDAP user. OME-Modular supports a maximum of two Active Directory or LDAP server configurations, each.

Topics:

- [Logging in to OME-Modular as local, Active Directory, or LDAP user](#)
- [OME-Modular home page](#)
- [Viewing device health](#)
- [Setting up chassis](#)
- [Initial configuration](#)
- [Configuring chassis settings](#)
- [Managing chassis](#)
- [Chassis groups](#)
- [Controlling chassis power](#)
- [Backing up chassis](#)
- [Restoring chassis](#)
- [Exporting chassis profiles](#)
- [Managing chassis failover](#)
- [Troubleshooting in chassis](#)
- [Blinking LEDs](#)
- [Interfaces to access OME-Modular](#)
- [Viewing chassis hardware](#)
- [Viewing chassis alerts](#)
- [Viewing chassis hardware logs](#)
- [Configuring OME-Modular](#)

Logging in to OME-Modular as local, Active Directory, or LDAP user

OME-Modular allows authentication for 64 local user accounts.

For Active Directory and generic LDAP user accounts, OME-Modular allows a minimum of one user account in a simple environment and a maximum of two accounts in a complex environment.

LDAP users can perform the following tasks using OME-Modular:

- Enable LDAP access
- Upload and view a Directory Service CA certificate
- Specify attributes while configuring LDAP. The attributes are—LDAP server address, LDAP server port, Bind DN, Bind password, user login attribute, group membership attribute, and search filter
- Associate an LDAP group with an existing or new management module role group

To log in as a local, Active Directory, or LDAP user:

1. Enter the **Username**.
2. Enter the **Password**.
3. Click **Login**.

After logging in successfully, you can do the following:

- Configure your account
- Change the password
- Recover the root password

Logging in to OME-Modular as Active Directory or LDAP user

To log in to OME-Modular as an Active Directory (AD) or LDAP user:

1. Add directory service
2. Import directory group
3. Log in with directory user credentials

To add directory service:

1. From the menu bar in the OME-Modular web interface, click **Application Settings > Users > Directory Services > Add**. The **Connect to Directory Service** window is displayed.
2. Select AD or LDAP, and enter the appropriate information.
3. If the directory type is AD, and the **Domain Controller Lookup** type is DNS, enter the domain name and group domain.

In the group domain, you can look for directory groups. You can include the directory groups as application users. You can also use the group domain for authenticating users during login. The format of the group domain can be—<Domain>.<Sub-Domain> or ou=org, dc=example, dc=com.

Use the "DNS" **Domain Controller Lookup** type, if you do not know the details of the domain controllers from which you want to import the group or groups. To use the DNS domain controller, ensure that you have done the following tasks on the **Network Settings** page:

- Selected the **Register with DNS** check box
- Provided the Primary and Alternate DNS server addresses

After you enter the domain name, OME-Modular searches the SRV records on the DNS servers to fetch the details of the domain controllers in that domain.

If you know the IP address or FQDN of the domain controllers, you can use the "Manual" **Domain Controller Lookup** type.

The **Test Connection** feature is only applicable to the "DNS" domain controller type.

Importing directory group

To import a directory group:

1. From the menu bar in the OME-Modular web interface, click **Application Settings > Users > Import Directory Group**. The **Import Directory** window is displayed.
2. Select the directory service from which you want to import the group.
3. Under **Available Groups**, select the group and click >>. The selected group is displayed under **Groups to be Imported**.
4. Assign a role to the imported groups.



You can import groups after assigning roles to them. A message is displayed after the groups are imported successfully. Users in the imported groups can access OME-Modular, with specific roles and privileges.

Logging in to OME-Modular using the directory user credentials

To log in to OME-Modular using the directory user credentials:

From the OME-Modular login page, log in using the AD user credentials. Enter the domain name, if necessary.

OME-Modular home page

When you log in to OME-Modular, the home page is displayed. This page displays a dashboard with high-level information about the system and the subcomponents. Use the search field on the page to search for settings available in OME-Modular. You can also view the job activity and events. To view the job activity, click  and to view events, click .

To return to the OME-Modular home page, click the OME-Modular logo or click **Home**.

- Chassis graphical view—On left of the page, a graphical view of the front and rear chassis is displayed. It shows all the modules (sleds, fans, power supplies, IOMs, and MMs) present in the chassis. A mouse over on each module displays a brief description and health status of the module. Click **View Devices** to see more details about the modules present in the chassis. Click **View Slot Information** to switch the display of the widget to slot information list.

- Slot information view—On the top left corner of the page, a list of modules present on the chassis is displayed showing slot information, health status and a link that goes into details. Modules in this list include compute, storage sleds, and IOMs. Click **View Inventory** to see more details about the modules present in the chassis. Click **View Chassis Image** to switch the display of the widget to chassis graphical view.
- **Chassis Information**—On the lower left corner of the page, you can view a summary of the chassis information such as service tag, asset tag, firmware version and power state.
- **Device Health**—On the upper right corner of the page, you can view the health status of chassis subsystems such as fans, power supplies, temperature and compute, networking, and storage sleds. When the subsystem status is unhealthy, you can click in the **Reason** to view the list of fault messages.
- **Recent Alerts**—On the top center of the page, you can view the most recent alerts for events occurring in the chassis. Click **View All**, to see all the alerts in the **Alerts** page.
- **Recent Activity**—Below the **Recent Alerts** widget, you can the most recent activities occurring in the chassis. Click **View All**, to view all the activities or jobs in the **Jobs** page.

NOTE: When you refresh inventory and power the chassis on after the chassis is AC power cycled, the inventory of the compute sled and IOM may be displayed after 3-5 minutes.

NOTE: If chassis has not been powered on after the AC Power Cycle operation, the inventory status is displayed as "unknown".

Viewing alerts

The **Alerts** section displays the specific types of alerts such as Critical, Warning, and Unknown. You can also view alerts for specific device types such as chassis, compute, networking, and storage.

Viewing jobs and activities

The **Recent Activity** section displays a list of recent jobs and activities, and their status. Click **All Activity** to go to the **Jobs** page and view detailed information about the jobs.

Multi-chassis management dashboard

Multiple chassis are grouped to form domains called Multi-Chassis Management (MCM) groups. An MCM group can have 20 chassis, where one is the lead and the remaining 19 are members. OME-Modular supports wired MCM groups where the chassis are daisy-chained through a redundant port on the management controller.

In a multi-chassis management (MCM) group, the number of events and jobs for the entire group is displayed. The **Device Health**, **Alerts**, and **Recent Activity** sections display the consolidated details of all the devices in the group.

NOTE: Maintain a minimum interval of two minutes between removing and inserting each device.

Viewing MCM home page

You can view the following information about the MCM group:

- MCM group—You can view:
 - Name of the group
 - Topology of the group using **View Topology**
 - Name, IP address, and service tag of the lead chassis
 - Name, IP address, and service tag of the member chassis
- **Device Health**—Displays the health status of the chassis subsystems—chassis, compute sled, networking, and storage. You can click the health status of the individual devices or click **All Devices**, to view a summary of the devices in the **All Devices** page.
- **Recent Alerts**—Displays the most recent alerts for events occurring in the lead chassis and the subsystems. Click **All Alerts**, to view the **Alerts** page for the lead and member chassis.
- **Recent Activity**—Displays the most recent activities occurring in the lead chassis and the subsystems. Click **All Activity**, to view the **Jobs** page for the lead and member chassis.

NOTE: If a member chassis is added to a chassis group based on a "Join Group" request from the member chassis, the status of the member chassis is displayed as "Unknown" for sometime, on the MCM dashboard.

Viewing lists of chassis in an MCM group

On the OME–Modular home page, the list of chassis that are part of the group is displayed on the left. The list displays the model, IP address, and the Service Tag of the chassis. The lead chassis is labeled for easy identification. Click the name of the chassis to access the details specific to the chassis. You can also use the listed IP address to directly access the OME–Modular web interface of the chassis.

Viewing device health

The **Devices > All Devices** page displays the health summary of the chassis, compute and storage sleds, and networking components.

A list of all the devices at the bottom of the **All Devices** page. You can select a device to view its summary on the right side of the list. You can sort the list using **Advanced Filters**.

You can also perform the following tasks on the **All Devices** page:

- Power control
- Update firmware
- Blink LED
- Refresh inventory

- NOTE:** When you initiate a Leave chassis group request while the inventory refresh is in-progress, an error message is displayed on the All Devices page even if the Leave Chassis Group task is successful.
- NOTE:** When a compute sled is inserted into a chassis, sometimes the message, "No device image found", is displayed. To resolve the issue refresh the inventory of the compute sled, manually.
- NOTE:** When you refresh inventory and power the chassis on after the chassis is AC power cycled, the inventory of the compute sled and IOM may be displayed after 3-5 minutes.

Setting up chassis

When you log in to the OME–Modular web interface for the first time, the configuration wizard is displayed. If you close the wizard, you can access it again by clicking **Configure > Initial Configuration**. This option is displayed only if the chassis is not yet configured.

To configure the chassis:

1. Log into OME–Modular.
The **Home** page is displayed.
2. Click **Configure > Initial Configuration**.
The **Chassis Deployment Wizard** is displayed.
For further steps, see [Initial configuration](#).

Initial configuration

Dell EMC recommends the following configuration threshold for better performance of the chassis. If the configuration exceeds the threshold, then some features including firmware update, backup, and restore may not work as expected. This may also affect system performance.

Component	Count
Templates	320
Alert Policy	50
Identity pool	501
Network (VLAN)	214
Catalog	50
Baseline	50

To configure a chassis:

1. Click **Devices > Chassis > View Details > Configure > Initial Configuration**.
The **Chassis Deployment Wizard** is displayed.

NOTE: You can configure the chassis using an existing chassis profile.

- In the **Import Profile** tab, click **Import** to open the **Import Profile** window. Enter details of the network share, where the chassis profile is located and click **Import**.
- In the **Time Configuration** tab, select the **Configure Time Settings** to configure the time zone and timestamp of the configuration.
- Select the **Use NTP** check box to configure the primary, secondary, or tertiary NTP addresses and click **Next**.

NOTE: It is recommended that at least three valid NTP servers, which synchronize to a single time source, are used to ensure reliable synchronization.

If you select multiple NTP servers, OME–Modular selects the NTP server algorithmically.

The **Activity and Alerts** tab is displayed.

- Configure the email, SNMP, and system log settings and click **Next**. The **iDRAC** tab is displayed.
- Select the **Configure iDRAC Quick Deploy Settings** check box to configure the password to access the iDRAC web interface and the management IPs, and click **Next**.
You can select the slots to which the iDRAC Quick Deploy settings must be applied.
The **Network IOM** tab is displayed.
- Select the **Configure I/O Module Quick Deploy Settings** check box to configure the password to access the IOM console and management IPs, and click **Next**.
The **Firmware** tab is displayed.
- Select the **Configure all devices to use following catalog** check box, select the network share type and click **Catalog** to open the **Add Firmware Catalog** window.
- Enter a name for the catalog, select the catalog source, and click **Finish** to save the changes and return to the **Chassis Deployment Wizard**.
- Click **Next** to view the **Proxy** tab and configure the proxy settings.
OME–Modular uses the proxy settings to access the Dell EMC website for the latest catalogs. You can also enable the HTTP proxy settings and proxy authentication.
- Click **Next** to view the **Group Definition** tab.
- Select **Create Group** to configure the chassis group settings.
- Click **Next** to view the **Summary** tab.

NOTE: After setting the time in the lead chassis, wait for the lead chassis time and the member chassis time to synchronize before performing any operation. The time configuration can be disruptive.

Configuring chassis settings

You can configure the following settings for a chassis:

- Power
- Network
- Network Services
- Local Access Configuration
- Location
- Quick Deploy

Configure chassis power

To configure the chassis power settings:

- Click **Devices > Chassis > View Details > Settings > Power**.
The **Power** configuration section is expanded.
- Select **Enable Power Cap** to specify the maximum power consumption capacity for the chassis. The **Power Cap** limits the power consumption of the chassis. When the power cap is reached, the sleds are throttled based on their power priority. You can specify the capacity in Watts, BTU/h, or percentage. The **Power Cap** option is displayed only if the **Enable Power Cap** check box is selected. The recommended power cap is 0-32767 Watts or 0-100 %. If you change the power cap in BTU/h, the power cap in W also changes. MX7000 chassis supports power sources of 110 Volts and 220 Volts.
- In the **Redundancy Configuration** section, select the required redundancy policy.

Power redundancy policies facilitate management of power consumption and power failure tolerance in the chassis. The available options are:

- **No Redundancy**—This policy distributes the enclosure power load across all PSUs. There are not any specific PSU population requirements for **No Redundancy**. The intent of the **No Redundancy** policy is to have the highest possible limit for power enablement of devices added to the enclosure. If there are single or multiple PSU failures, then the enclosure limits the performance to operate within the power capabilities of the remaining PSUs.
- **Grid Redundancy**—This policy distributes the enclosure power load across all PSUs. The six PSUs are organized into two groups: Grid A consists of PSUs 1, 2, 3, and Grid B consists of PSUs 4, 5, 6. It is recommended that the PSUs are populated in the following order: 1, 4, 2, 5, 3, 6, where an equal number of PSUs on each grid is optimized for Grid Redundancy. The grid with the largest PSU capacity determines the limit for power enablement of devices added to the enclosure. If there is a grid or PSU failure, then the enclosure power is distributed among the remaining PSUs with the intent that a single healthy grid continues to provide power to the system without degrading the performance.
- **PSU Redundancy**—This policy distributes the enclosure power load across all PSUs. There are no specific PSU population requirements for redundant PSUs. PSU redundancy is optimized for a population of six PSUs, and the enclosure limits the power enablement of devices to fit within five PSUs. If there is a single PSU failure, then the enclosure power is distributed among the remaining PSUs without degrading the performance. If there are less than six PSUs, then the enclosure limits the power enablement of devices to fit within all populated PSUs. If there is a single PSU failure, then the enclosure limits the performance to operate within the power capabilities of the remaining PSUs.

4. In the **Hot Spare Configuration** section, select the **Enable Hot Spare** configure the Hot Spare primary grid.

The Hot Spare feature facilitates voltage regulation when power utilization by Power Supply Unit (PSUs) is low, considering the total output capacity of the PSU. By default, the Hot Spare is enabled. When the Hot Spare is enabled, a redundant PSU is put in sleep state when the power utilization is low. The Hot Spare is not enabled if the:

- PSU redundancy is inactive.
- Power budget of the system configuration exceeds the PSU output capacity.
- Grid Redundancy Policy is not selected.

The MX7000 PSUs support the Hot Spare feature with three PSU pairs. The feature enables a PSU pair to have one active PSU and one PSU in sleep mode while the enclosure power consumption is low, and the three PSU pairs meet all the power requirements for the enclosure. This enables efficient power utilization when the overall enclosure power requirement is low. The partner PSU wakes the paired PSU from sleep mode by sending a WAKE signal when the enclosure power requirement increases. The PSU pairs for MX7000 are PSUs: 1 & 4, 2 & 5, and 3 & 6.

5. From the **Primary Grid** option, select the PSU where you want to enable the Hot Spare, from the drop-down.

6. Click **Apply** to save the chassis power settings.

Configure chassis management network

You can configure the network settings for the management modules that are inserted into an MX7000 chassis.

- LAN/NIC interface
- IPv4
- IPv6
- DNS Information
- Management VLAN

To configure the chassis network:

1. Click **Devices > Chassis > View Details > Settings > Network**.

The **Network** configuration section is expanded.

2. In the **General Settings** section, you can enable or disable NIC, **Register with DNS**, and **Auto Negotiation**. By default, the **Enable NIC** check box is selected.

If you enable **Register with DNS**, then enter the **DNS Name** of the chassis that you want to register with a DNS server. You can access OME-Modular using the existing FQDN even after the **Register with DNS** option is disabled in the application. This is because the earlier option remains in the Network cache or the DNS Server cache, based on the configured Time to live (TTL).

NOTE: You can only access the FQDN temporarily.

NOTE: Clear the cache in the DNS after the Register with DNS is disabled, to prevent logging in with the FQDN address.

NOTE: If the Register with DNS option is enabled, you cannot modify the Enable VLAN option.

3. Enter the **DNS Name**. The DNS name can have a maximum of 58 characters. The first character must be an alphanumeric character (a-z, A-Z, 0-9), followed by numeric characters or a hyphen (-).
4. Enable or disable the **Use DHCP for DNS Domain Name** option and turn the **Auto Negotiation** on or off.

If the **Use DHCP for DNS Domain Name** is disabled, then enter the **DNS Domain Name**.

NOTE: You can enable Use DHCP for DNS Name only if IPv4 or IPv6 has DHCP configured. OME–Modular obtains its DNS domain name from either a DHCP or DHCPv6 server when Use DHCP for DNS Name is enabled.

If **Auto Negotiation** is false or disabled, you can choose network port speed.

NOTE: Setting Auto Negotiation to false and choosing a network port speed may result in the chassis losing link to the network switch in Top of Rack, or to the neighbor chassis, if running MCM. It is recommended that the Auto Negotiation is set to true for most use cases.

Table 2. Top of the Rack Support Matrix for management module and management module uplink

Top of the Rack Switch Configuration	Management Module Configuration	Supported for Management Module Uplink(YES/NO)
100 Mbps (Auto negotiation OFF)	100 Mbps (Auto negotiation OFF)	YES
10 Mbps (Auto negotiation OFF)	10 Mbps (Auto negotiation OFF)	YES
Auto Neg ON	Auto Negotiation ON	YES
100 Mbps (Auto negotiation OFF)	Auto Negotiation ON	NO
10 Mbps (Auto negotiation OFF)	Auto Negotiation ON	NO
Auto Negotiation ON	100 Mbps (Auto negotiation OFF)	NO
Auto Negotiation ON	10 Mbps (Auto negotiation OFF)	NO

5. In the **IPv4 Settings** section, configure the following:

- **Enable IPv4**
- **Enable DHCP**
- **IP Address**
- **Subnet Mask**
- **Gateway**
- **Use DHCP to Obtain DNS Server Addresses**
- **Static Preferred DNS Server**
- **Static Alternate DNS Server**

6. In the **IPv6 Settings** section, configure the following:

- **Enable IPv6**
- **Enable Autoconfiguration**
- **IPv6 Address**
- **Prefix Length**
- **Gateway**
- **Use DHCPv6 to Obtain DNS Server Addresses**
- **Static Preferred DNS Server**
- **Static Alternate DNS Server**

NOTE: The static IPv6 IP address that is already configured is applied and displayed in OME–Modular when the configuration is changed from static to DHCP IP.

7. Enable or disable the VLAN for the chassis. You can configure the VLAN settings only if the **Register with DNS** check box is cleared. You can change from a VLAN network to a non-VLAN network, or move from a non-VLAN network to a VLAN network, only if **Register with DNS** check box is cleared.

By default, the IPv4 settings are enabled and the DNS registration is disabled with a default name. You can modify the name using any local interfaces such as OpenManage Mobile.

NOTE: Ensure that the network cable is plugged to the correct port when you modify the VLAN state for the change to be effective.

Isolate the chassis management from the data network as the uptime of a chassis that is improperly integrated into your environment cannot be supported or guaranteed. Due to the potential of traffic on the data network, the management interfaces on the internal management network can be saturated by traffic that is intended for servers. This results in OME-Modular and iDRAC communication delays. These delays may cause unpredictable chassis behavior, such as OME-Modular displaying iDRAC as offline even when it is up and running, which in turn causes other unwanted behavior. If physically isolating the management network is impractical, the other option is to separate OME-Modular and iDRAC traffic to a separate VLAN. OME-Modular and individual iDRAC network interfaces can be configured to use a VLAN.

NOTE: Any change in the attribute settings leads to IP drop or unavailability of the OME-Modular web interface for some time. However, the OME-Modular web interface recovers automatically.

8. Click **Apply** to save the chassis network settings.

Configure chassis network services

The chassis network services configuration consists of SNMP, SSH, and remote RACADM settings.

To configure network services:

1. Click **Devices > Chassis > View Details > Settings > Network Services**.

The **Network Services** section is expanded.

2. In the **SNMP Settings** section, select the **Enabled** check box to enable the SNMP settings and select the **Port Number**.

The port number can be 10-65535.

NOTE: For SNMP operations, configure the timeout parameter on the client to facilitate successful completion of the task. You may have to adjust the timeout parameter based on the network latency.

3. Enter the SNMP **Community Name**. The length of the community name must be less than or equal to 32 characters.
4. Download the **Management Information Base (MIB) file** to a local drive on your system.
5. In the **SSH Settings** section, select the **Enabled** check box to enable the SSH settings for the chassis and select the maximum number of SSH sessions.

By default, a chassis can have a maximum number of four SSH sessions.

6. Select the **Idle Timeout**, in seconds, for which the SSH session can remain idle. The session is terminated when the timeout is reached. The valid range is 60-10800 seconds. The default value is 1800 seconds. Select 0 seconds to disable the timeout.
7. Select the SSH **Port Number**. The port number can be 10-65535.

The default port number is 22.

8. Enable the remote RACADM session for the chassis.

You can view the remote RACADM option on the web interface only if you have the chassis administrator privilege.

NOTE: A log for remote RACADM session (login or logout) is displayed in the Audit Logs page, irrespective of the remote RACADM status. If the remote RACADM option is disabled, the feature does not work.

NOTE: Any change in the attribute settings leads to IP drop or unavailability of the OME-Modular web interface for some time. However, the OME-Modular web interface recovers automatically.

9. Click **Apply** to save the chassis network services settings.

Configure local access

You can configure chassis power button, quick sync, KVM, LCD, and chassis USB direct accesses for a chassis.

To configure the local access settings in a chassis:

1. Click **Devices > Chassis > View Details > Settings > Local Access Configuration**.

The **Local Access Configuration** section is expanded.

2. Select **Enable Chassis Power Button** to use the power button to turn the chassis off or on.

If the check box is cleared, you cannot change the power state of the chassis using the chassis power button.

3. Select the **Quick Sync access** type.

The available options are:

- Read-only—Enables read-only access to WiFi and Bluetooth Low Energy (BLE). You cannot write configuration information using quick sync.
- Read-write—Enables writing configuration using quick sync.
- Disabled—Disables reading or writing configuration through quick sync.

NOTE: The Quick Sync feature uses a lower radio frequency (RF) power when advertising and increases the RF power after the certificate authentication. The RF range is based on the environment and can vary.

4. Select **Enable Inactivity Timeout** to enable the idle timeout and enter the **Timeout Limit**.

Timeout is the idle time when there is no Wi-Fi traffic. Specify inactivity timeout limit, in seconds. The timeout can be between two minutes and 60 minutes.

NOTE: The Timeout Limit option is available only if the Enable Inactivity Timeout is selected.

5. Select **Enable Read Authentication** to use your user credentials to read the inventory in a secure data center. By default, this option is selected. If you clear this check box, you cannot access the secure data center.
6. Select **Enable Quick Sync Wi-Fi** to use WiFi to communicate with the chassis. By default, the **Enable Quick Sync Wi-Fi** check box is selected.
7. Select **Enable KVM Access** to configure the quick sync setting using KVM. You can also use the RACADM or Redfish command to enable or disable KVM. For more information, see the *OME - Modular for PowerEdge MX7000 Chassis RACADM CLI Guide* available at www.dell.com/openmanagemanuals.

You can use the DisplayPort in the chassis to stream the video in the KVM. If the external DP to Video Graphics Array (VGA) converter is available, you can stream the KVM video in the VGA too.

8. Select the **LCD Access** option for quick sync.

The available options are:

- Disabled
- View Only
- View and Modify

NOTE: The LCD Access option is displayed only if there is a system with LCD available in the chassis.

9. In the **User Defined** text box, enter the text that you want to see on the LCD Home screen. The LCD Home screen is displayed when the system is reset to factory default settings. The text can have a maximum of 62 characters and supports a limited number of UTF-8 characters. If a UTF-8 character that is used in the text is not supported, a box is displayed instead of the character. The default string is the service tag of the system.
10. From the **LCD Language** drop-down, select the language in which the text on the LCD must be displayed.

The available options are:

- English
- French
- Spanish
- German
- Japanese
- Chinese

By default, the text is displayed in English.

11. Select the **Enable Chassis Direct Access** text box to enable accessing the MX7000 chassis from a host such as a laptop or server, using a USB On-The-Go (OTG) cable.

If the **Enable Chassis Direct Access** check box is cleared, the existing chassis direct sessions are disconnected and the Chassis Direct LED turns off. When the feature is disabled, you cannot connect the laptop to the chassis. The URL `https://ome-m.local` is inaccessible. After enabling the feature, reattach the USB cable and wait for Chassis Direct LED to turn green to access the chassis phonebook. For more information, see the section, [Chassis Direct](#).

12. Click **Apply** to save the quick sync settings.

Chassis Direct

The Chassis Direct feature in OME-Modular enables users to access management consoles such as iDRAC and management module of devices on the chassis. The MX7000 chassis has several USB ports. The Right Control Panel (RCP) on the front of the chassis has three USB ports. Two ports are regular sized USB-A ports, for keyboards and mouse used for the chassis level KVM. The third port is a Micro-AB port that supports USB OTG. To use Chassis Direct, connect the USB OTG port to a laptop. The Poleg processor on the management module emulates a USB network interface and provides a network bridge into the management VLAN. The network is same that QuickSync 2 bridges for OpenManage Mobile WiFi access.

Remove the USB cable that is connected to the front panel and AC power cycle the chassis.

With the laptop that is connected to the USB OTG port on the chassis, you can access the MM user interface and iDRAC user interface or KVM. You can get access by launching a browser on the laptop and entering the URL, `https://ome-m.local`. A chassis phonebook page which contains a list of entries to the available devices on the chassis, is displayed. This option provides a better experience than the front panel KVM, which provides access only the command-line prompt access for OME-Modular.

Select the check box to enable accessing the MX7000 chassis from a host such as a laptop or server, using a USB On-The-Go (OTG) cable. Connect the USB OTG cable from the host to the micro USB port on the front panel (right control panel) of the MX7000 chassis. On successful connection, LED under the micro-USB on the right control panel of MX7000 chassis turns green and USB Ethernet adapter is displayed on the host. The chassis is automatically configured with an IPV4 and IPV6 address. After ensuring that the addresses are configured, open a web browser and enter the URL, `https://ome-m.local` in the address bar.

When you enable or disable the Chassis Direct feature in OME-Modular, the following error codes are displayed:

The Chassis Direct feature in OME-Modular has a mutual exclusivity with the Quick Sync feature. Before downgrading the management module firmware from 1.10.00 version to an earlier version, remove the USB cable that is connected to the front panel of chassis. If USB cable is not removed and 1.10.00 firmware is downgraded, the Quick Sync feature may be degraded. AC power cycle of chassis to restore Quick Sync back to health.

- The chassis has Quick Sync and the Chassis Direct feature is enabled. That is, the USB cable is attached to the USB connector on the front panel.
- The management module version is downgraded from 1.10.00 to an earlier version.

Table 3. Chassis Direct—LED blink status and description

Error code	Chassis Direct LED blink status	Description and resolution
1	Amber	The USB network link is down as the Chassis Direct feature is disabled. Resolution —Enable the Chassis Direct and reattach the USB cable to access the chassis phonebook.
2	Amber	The USB network link does not come up as the chassis internal USB operation failed. Resolution —If the issue persists, reattach the USB cable to the laptop or perform and AC power cycle of the chassis.
3	Amber	The USB network link fails to come up owing to an issue on the host laptop. Resolution —If the issue persists, reattach the USB cable.
4	Turned off	The USB network link is down as the USB cable is disconnected. Resolution —Reattach the USB cable for the link to come up.

If the Chassis Direct feature is disabled and the USB cable is inserted, the Chassis Direct LED turns amber and the alert, USR0197, is displayed on the OME-Modular web interface. You can see the alert only if you have logged in to OME-Modular using the public network. If you repeat the action within a short interval, the alert is not displayed. However, the Chassis Direct LED remains amber as the MM suppresses consecutive duplicate alerts.

Configure chassis location

To configure the location of the chassis:

1. Click **Devices > Chassis > View Details > Settings > Location**.
The **Location** configuration section is expanded.
2. Enter the location names for the **Data Center**, **Room**, **Aisle**, and **Rack**.
The **Data Center**, **Room**, **Aisle**, and **Rack** support up to 128 characters.
3. Enter the number of the **Rack Slot** and the name of the **Location** where the rack is located.
The **Rack Slot** supports 1-255 numeric characters.

The **Location** supports up to 128 characters. It is supported for backward compatibility. The Data Center, Aisle, Rack, and Rack Slot properties replace this property. Use these properties to describe the physical location of the chassis.

4. Click **Apply** to save the location settings.

Configure Quick Deploy settings

The **Quick Deploy** feature enables you to configure the password to access the iDRAC user interface, IOMs, and IPv4 and IPv6 settings. These settings can be applied to existing compute sleds or IOM devices immediately. You can apply the **Quick Deploy** settings to compute sleds when they are inserted into the chassis, later. However, you cannot apply the **Quick Deploy** settings to IOMs that are inserted later.

Quick deploy settings are validated when the job is run. If an invalid parameter is used, the quick deploy job fails. The **Quick Deploy** job parameters are not evaluated, as they can contain any value, which is delegated while running the job.

Enabling and disabling quick deploy is a web interface feature to determine if the controls are enabled to configure **Quick Deploy** settings. The back-end only processes requests from the web interface.

NOTE: After the quick deploy settings are applied to the compute sled, the IP configuration is displayed in the OME-Modular web interface, when the inventory is refreshed.

NOTE: When IPv4 for IPv6 is disabled for FC IOMs, the Device IPv4 address or Device IPv6 address is blank on the Quick Deploy page for IOMs. However, for network IOMs, the IPv4 and IPv6 device addresses are, :: and 0.0.0.0.

To configure the **Quick Deploy** settings:

1. Click **Devices > Chassis > View Details > Settings > Quick Deploy**.

The **Quick Deploy** configuration section is expanded.

2. Enter and confirm the password to access the iDRAC user interface.

The password can be up to 20 characters in length.

NOTE: If any iDRAC IP configuration is modified, the SSO for the SLEDs is functional from the OME-Modular console only after the default inventory task or manual inventory refresh is complete.

3. In the **Management IP** section, select **IPv4 Enabled** to enable the IPv4 network settings and select the **IPv4 Network Type**.

The available options are:

- Static
- DHCP

4. Enter the **IPv4 Subnet Mask** and **IPv4 Gateway**.

NOTE: The **IPv4 Subnet Mask** and **IPv4 Gateway** options are displayed only if the **IPv4 Network Type** is "Static".

5. Select **IPv6 Enabled** to enable the IPv6 network settings and select the **IPv6 Network Type**.

The available options are:

- Static
- DHCP

6. If the **IPv6 Network Type** is Static, select the **IPv6 Prefix Length** and enter the **IPv6 Gateway**.

7. From the list of slots that is displayed, select the check box next to the slot number to which you want to apply the **Quick Deploy** settings.

8. In the **Network IOM Settings** section, enter and confirm the password to log in to the IOM interface.

9. Select **IPv4 Enabled** to enable the IPv4 network settings and select the **IPv4 Network Type**.

The available options are:

- Static
- DHCP

10. Enter the **IPv4 Subnet Mask** and **IPv4 Gateway**.

NOTE: The **IPv4 Subnet Mask** and **IPv4 Gateway** options are displayed only if the **IPv4 Network Type** is "Static".

11. Select **IPv6 Enabled** to enable the IPv6 network settings and select the **IPv6 Network Type**.

The available options are:

- Static

- DHCP

12. If the **IPv6 Network Type** is Static, select the **IPv6 Prefix Length** and enter the **IPv6 Gateway**.
13. Click **Apply** to save the **Quick Deploy** settings.

Managing chassis

You can view the list of chassis and the chassis details on the **Chassis** page. The details are—health, power state, name, IP address, service tag, and model of the chassis. You can also select a chassis to view the graphical representation and summary of the chassis, on the right side of the **Chassis** page.

You can also perform the following tasks on the **Chassis** page:

- Control chassis power
- Update firmware
- Blink LED
- Refresh chassis inventory
- Filter the chassis list

i **NOTE:** When a chassis is power cycled, the inventory of the compute sleds and IOMs may be displayed in the OME-Modular web interface after three to five minutes.

i **NOTE:** Maintain a minimum interval of two minutes between removing and inserting each device.

i **NOTE:** After a chassis power off, the compute SLEDs are polled based on the event from the chassis. Each event from the chassis triggers a health-poll. You may see multiple connection loss events from compute SLEDs.

Creating chassis filters

You can sort the list of chassis that are displayed on the **Devices > Chassis** page, using filters.

To create filters:

On the **Chassis** page, click **Advanced Filters** to view the filter options. The following options are displayed:

- **Health**
- **State**
- **Name Contains**
- **IP Address Contains**
- **Service Tag Contains**
- **Model**

Viewing chassis overview

On the chassis **Overview** page, you can click **View Slot Information** to view the compute sled slot details. A graphical representation of the chassis is displayed on the left side. Information about the chassis is displayed below the graphical representation. The information includes FIPS status of the chassis, name, model, service tag, asset tag, express service code, management IP, firmware version, power state, and faceplate power of the chassis. Click **View Devices** to view the list of all devices on the **All Devices** page.

You can also see information under the following sections:

- **Chassis Subsystems**—Displays the health status of the chassis components such as battery, fan, IOMs, and power supply. Fabric Consistency Check (FCC) information and health change is displayed under **Chassis Subsystems**. But the FCC details of the compute sled are not displayed in the chassis graphical representation and the compute **Overview** page.
- **Environment**—Displays the power consumption units and temperature of the chassis. Click **View Power Statistics** to view the chassis power consumption details such as current redundancy state, peak headroom, and system energy consumption. Click **Power Usage** to view the chassis power supply information on the **Chassis > Hardware > Chassis Power Supplies** page. If a failover or management module reboot is performed, the last reset power statistics timestamp is updated based on the failover or management module reboot timestamp.

i **NOTE:** The temperature statistics timestamp remains unchanged after a failover or management module reboot.

- **Recent Alerts**—Displays the number and details of the tasks that are performed in the chassis. Click **View All** to view the list of all alerts that are related to the compute sled on the **Chassis > Alerts** page.
- **Recent Activity**—Displays the status of the jobs that are performed in the compute sled.
- **Server Subsystems**—Displays a summary of information about the server sub systems. The information includes the health status of the components such as battery, memory, processor, and voltage.

If you have the Chassis Administrator privileges, you can perform the following tasks in this tab:


- **Power Control** tasks:
 - **Power Off (Non-graceful)**—Turns off the server power, which is equivalent to pressing the power button when the server is turned on. This option is disabled if the server is already turned off. It does not notify the server operating system.
 - **Power Cycle System (Cold Boot)**—Turns off and then restarts the server (cold boot). This option is disabled if the server is already turned off.
 - ⓘ **NOTE: When the chassis is power cycled all devices in the chassis are also powered cycled. The management module does not get power cycled. But, alerts are logged indicating that the connectivity to devices is lost owing to the power cycle operation.**
 - **Power Off (Graceful)**—Notifies the server operating system to turn off the server. This option is disabled if the server is already turned off.
- Configuration tasks:
 - **Create Chassis Group**
 - **Join Chassis Group**
 - **Initial Configuration**
- Troubleshooting tasks:
 - Extract Log—You can extract the logs to a CIFS or NFS share, or a local drive on your system.
 - Diagnostic Commands
 - Reset management module
 - Terminate serial connection
- Turn-on or turn off LEDs using **Blink LED**.
- Back up, restore, export chassis profile, and perform failover.

ⓘ **NOTE: After a chassis power off, the compute SLEDs are polled based on the event from the chassis. Each event from the chassis triggers a health-poll. You may see multiple connection loss events from compute SLEDs.**

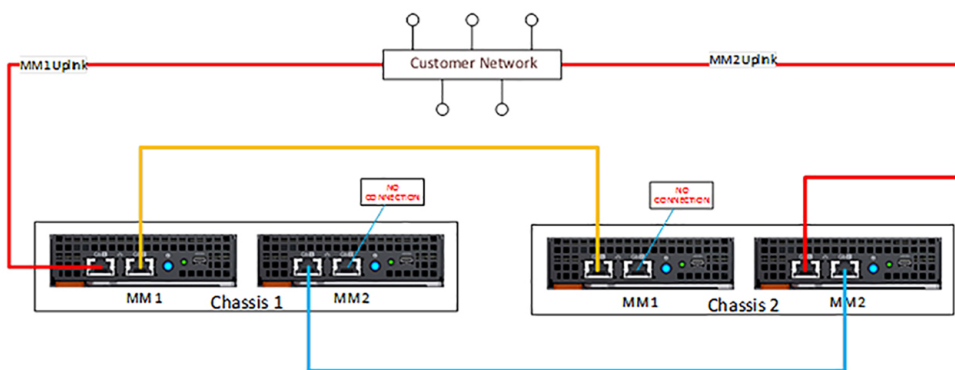
Wiring chassis

The automatic uplink detection and network loop prevention features in OME-Modular facilitate connection of multiple chassis with cables. The wiring saves port usage in the data center switches and access each chassis in the network. The cabling or wiring of chassis in this way is called stack.

While wiring a chassis, connect one network cable from each management module to the Top of Rack (ToR) switch of the data center. Ensure that both the ports on the ToR are enabled and are in the same network and VLAN. The following image is a representation of the individual chassis wiring:



The following image is a representation of the two-chassis wiring:



Chassis groups

You can group many chassis to form a multi-chassis management (MCM) group. An MCM group can have one lead chassis and 19 member chassis. You can use any management module to create an MCM group. The management module that is used for creating the MCM is the leader of the group, by default. The MCM group is of wired type, where the chassis is daisy-chained or wired through a redundant port on the management module. The chassis that you select for creating the group must be daisy-chained to at least one chassis. You can view a list of wired chassis and select all or the required number of chassis for creating the MCM group.

NOTE: You must have the chassis administrator privilege to create an MCM group.

You can perform the following tasks using an MCM group:

- View the health of the MCM group and the member chassis.
- Automatically apply settings of the leader chassis to member chassis.
- Perform any chassis operation on the MCM group.

You can add member chassis to an MCM group in two ways:

- Automatic—Enables automatic inclusion of the member to the chassis group. The automatic inclusion process does not require approval from the chassis administrator.
- Manual—Mandates approval by the chassis administrator to include the member chassis to the chassis group.

Prerequisites for creating a wired group

Following are the prerequisites to create a wired or daisy-chained chassis group:

- List of wired daisy-chained chassis—All the chassis must be on the private stack. You need not enter a password as the machine to machine authentication trust is used.
- Ensure that you have added member chassis to the group using the automatic or manual method.
- Ensure that the chassis settings are selected for applying to the other chassis—Power, user authentication, alert destination, time, proxy, security, network services, local access.
- Ensure that Auto Negotiation is set to true in all the chassis that are connected to form an MCM group. For more information, see [Configuring chassis network](#).
- Before stacking the chassis for creating a group or adding new members to the existing group, ensure that all the chassis have the same OME-Modular firmware version.

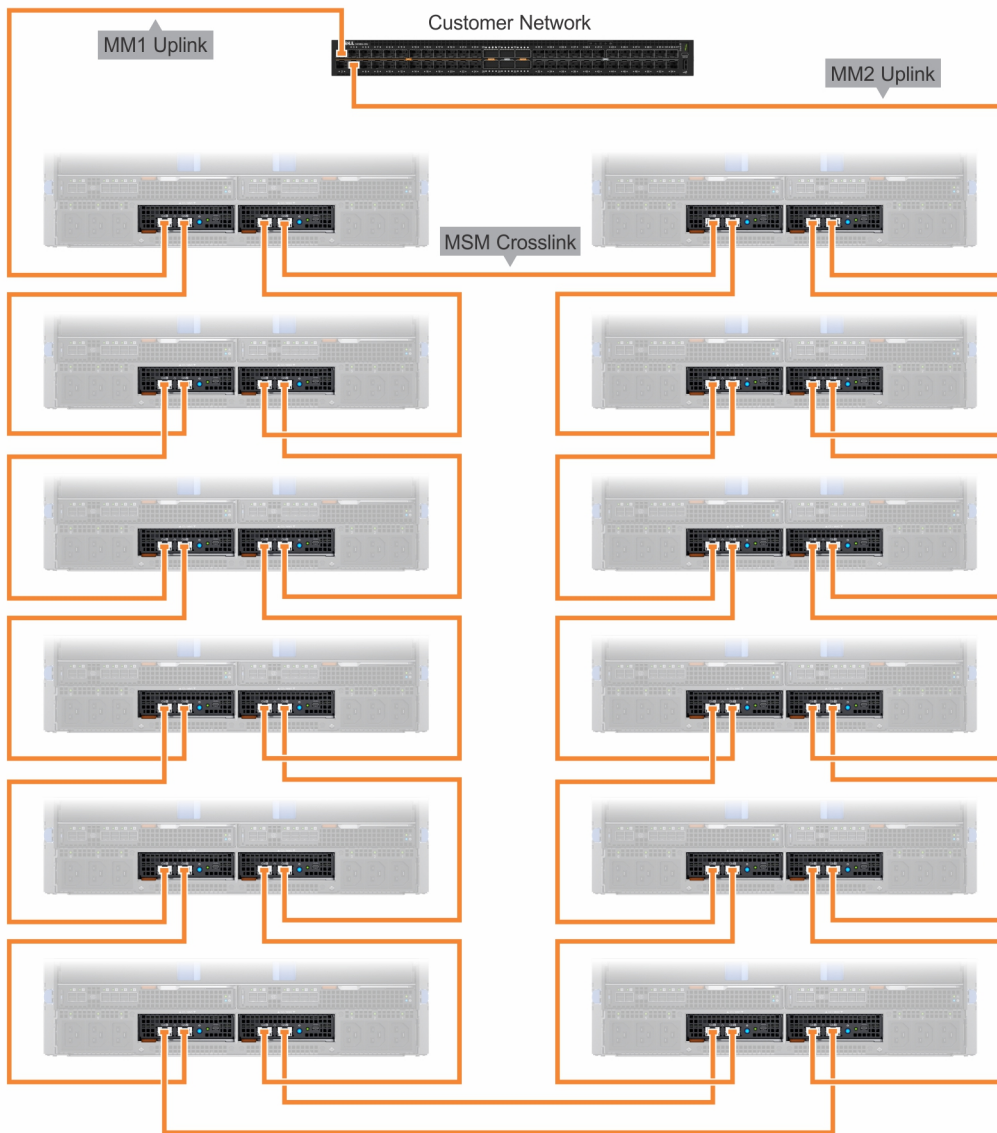
Before creating an MCM group, ensure that the MX7000 management networks are wired together in a stacked configuration. The stacked configuration helps in surviving:

- A single network cable failure
- A single management module failure
- Power loss owing to any chassis in the stack
- Failover of a chassis in the stack

NOTE: If any of the issues that are listed above occur, the management network access to all components in the daisy-chained group may be interrupted for up to 10 minutes. The OME - Modular web interface recovers automatically.

The wired chassis are displayed as under **Available Chassis** in the **Group Deployment Wizard**.

The following image is a representation of the recommended MCM wiring:



Creating chassis groups

To create a chassis group:

1. On the chassis dashboard, click **Overview > Configure > Create Chassis Group**.
The **Create a Group and Configure Lead Chassis** wizard is displayed.
2. Enter a name and description for the chassis group you want to create.
The group names can contain letters and numbers and must be fewer than 48 characters. However, the group names cannot contain spaces and special characters.
3. Select the onboarding permission type.
4. Select the configuration settings that you want to propagate to the member chassis.
The settings are:
 - All—Applies all settings of the lead chassis to the member chassis
 - Power—Cap, redundancy, compute sled priority
 - User Authentication—Directory services, local users
 - Alert Destination—Email, SNMP trap, system log
 - Time Settings—Date, time, time zone, NTP
 - Proxy Settings—All settings
 - Security Settings—Login IP range, log on lockout policy
 - Network Services—SNMP, SSH, remote racadm, web server

- Local Access Configuration—Chassis power button, quick sync, KVM, LCD, serial access

5. Click **Next** to view the summary of the group.

The dashboard of a leader chassis displays a summary of the health information, recent activity, and recent alerts of the member chassis. You can select a member chassis to view its details.

The current membership ID of the chassis is displayed on the left side.

Adding member chassis to groups

You can add members to the chassis groups from the **Overview** page of the lead chassis or from the member chassis.

Adding member chassis from lead chassis

To add a member chassis to the group from the lead chassis:

1. On the lead chassis **Overview** page, click **Configure > Add member**.
The **Add Chassis** window is displayed. The discovered chassis are displayed under **Available chassis**.
2. Select the number of chassis you want to add to the chassis group and click **Add**.
The list of added chassis is displayed at the bottom of the window.
3. Click **Finish**.

Adding individual chassis to chassis groups

To add an individual chassis to the chassis group:

1. On the chassis **Overview** page, click **Configure > Join Chassis Group**.
The **Join Group** window with all the existing MCM groups in the stack is displayed.
2. Select the chassis or MCM group to which to want to add the member, from the **Select a Group** drop-down.
3. Click **Finish**.

If the MCM group is created with manual on boarding policy, the join request is in the pending list for the lead chassis to confirm the addition of the member chassis. The lead chassis can approve or reject the request.

If the MCM group is created with automatic on boarding policy, no approval is required from the lead chassis. The individual chassis is automatically added to the MCM group to become a member chassis.

4. Log in to the lead chassis and approve the request of the member chassis to join the chassis group.

Assigning backup lead

In a multi-chassis environment, the lead chassis may sometimes fail temporarily or retire. In such situations, it is necessary to nominate a member chassis in the MCM group as a backup to the lead chassis. The backup lead chassis is promoted as a lead chassis when the existing lead chassis fails or retires.

1. On the MCM dashboard, click **Configure > Edit Backup Lead Settings**.
The **Edit Backup Lead Settings** window is displayed.
If a backup is already assigned, the name of the backup chassis is displayed in the **Current backup** field.
2. From the **Assign backup** drop-down, select the name of the member chassis that you want to select as the backup lead chassis.
3. Click the **Lead Virtual IP Configuration (Optional)** and select the **Enable Virtual IP** check box.
The virtual IP, if configured, facilitates consistency in the IP when the lead chassis role is transferred from one chassis to another.
4. Click **Additional Information** to view details about enabling the virtual IP. The details are:
 - **Modifying the network settings may impact the virtual IP configuration**
 - **Disabling the NIC will also disable the virtual IP**
 - **Disabling IPv4 will not disable the virtual IP**
 - **Enabling VLAN will leave the virtual IP accessible only within the specified VLAN**
 - **Enabling/disabling DHCP for IPv4 will reconfigure the virtual IP to match the new subnet mask and gateway**

Also, see the section, [Use case scenarios](#).

When a job for assigning a member chassis as the backup lead is stopped, the status of the job on the **Jobs** page is displayed as **Stopped**. However, the member chassis is assigned as the backup lead of the group.

Promoting backup chassis as lead

You can promote the backup chassis as the new lead chassis when the existing lead chassis fails. If the initial lead chassis is available, you can also assign it as a member chassis. To promote the backup chassis as the lead chassis, you must log in to the backup chassis.

After promoting a backup lead as lead chassis, detach and reattach any profiles that are attached to a slot containing a compute sled, in the new lead. Detaching and reattaching the profiles ensures that the assignment is persistent. The "promote" task does not affect profiles that are assigned to empty slots. Also, see the section, [Use case scenarios](#).

1. On home page of the backup chassis, click **Configure > Promote As Lead Chassis**.
The **Promote As Lead Chassis** window is displayed.
2. Click **Promote**.

After promoting the backup lead as the new lead of the chassis group, If you have the chassis administrator privileges, perform the following steps before putting the old lead chassis back into the production environment:

1. From the new lead chassis, remove the old lead chassis from the group to remove all references to the old lead chassis.
2. Remove the old lead chassis from the stacking network.
3. Run a forced reset configuration task by using the REST API, `/api/ApplicationService/Actions/ApplicationService.ResetApplication` For details, see the *OpenManage Enterprise and OpenManage Enterprise - Modular Edition RESTful API Guide*.

The reset configuration task transitions the old chassis to a stand-alone chassis and ready to be part of the production environment.

When a backup lead is promoted as the lead chassis, join requests from other member chassis sent to the earlier lead chassis, are not displayed on the MCM dashboard of the new lead. As a result, the particular member chassis cannot send joining requests to other groups in the stack. To unblock the pending requests, run the following API from the member chassis from which the joining requests were sent and resend the requests:

URI—`/api/ManagementDomainService/Actions/ManagementDomainService.DeletePendingDomains`

Method—POST

Payload—empty

Retiring lead chassis

You can use the retirement process of the existing lead chassis to make it a member chassis of the existing group or a stand-alone chassis.

1. On the MCM dashboard, click **Configure > Retire Lead Chassis**.
The **Retire Lead Chassis** window is displayed.
2. Select one of the following options:
 - Make it a member of the current group.
 - Make it a standalone chassis.
3. Click **Retire**.
Also, see the section, [Use case scenarios](#).

MCM dashboard

The MCM dashboard is displayed only when a multi-chassis management (MCM) group is created. You can view the name of MCM group on the left side of the dashboard. Below the group name, you can view the names, IPs, and service tags of the lead and member chassis. The lead chassis is indicated by the "LEAD" on the right side of the chassis name and the backup chassis is indicated by "BACKUP".

Click **View Topology** to view the structure of the MCM group.

The mid section of the MCM dashboard displays the health summary of all chassis, compute, networking, and storage devices in the MCM group. You can view the list of all the devices in the group by clicking **All Devices** at the upper right corner of the dashboard.

Below the health summary, you can view the alerts that are based on criticality of the alert and device type. Click **All Alerts** to view the list of the alerts that are related to all events in the MCM group.

You can view the details of the recent activities that are related to the group on the right side of the dashboard. The details comprise of the name and status of the activity, and timestamp of the activity. Click **All Activity** to view a list of all the activities that are related to the group, on the **Jobs** page.

Controlling chassis power

You can turn on and turn off the chassis power supply from the OME-Modular home page.

If you turn off the chassis manually or a power grid failure results in turning off multiple chassis, IOMs, and compute sleds, then, turning on all the chassis and compute sleds may result in failed inventory jobs for two to three hours. However, the inventory jobs recover with no impact to the chassis and related components.

To control the chassis power:

1. On the home page, click **Power Control** and select the required option.

The available options are:

- Power Off (Non-graceful)
- Power Cycle System (Cold Boot)
- Power Off (Graceful)

i **NOTE:** After login, wait for 7 minutes, if the IP is unavailable, then check if:

- **The cable is connected.**
- **DHCP is configured, ensure that the cable is connected to a Top of Rack (TOR) switch that has connectivity to the DHCP server.**

A message is displayed prompting you to confirm your action.

2. Click **Confirm** to proceed.

Backing up chassis

You must back up the chassis and compute sled configuration for later use. To backup the chassis, you must have administrator access with the device configuration privilege. The chassis configuration contains the following settings:

- Setup configuration
- Power configuration
- Chassis network configuration
- Local access configuration
- Location configuration
- Slot configuration
- OME–Modular network settings
- Users settings
- Security settings
- Alert settings

You can use the backed-up configuration in other chassis.

To create a chassis backup:

1. On the chassis **Overview** page, click **More Actions > Backup**.
The **Backup Chassis** window is displayed.
2. In **Backup File Location**, select the **Share Type** where you want to store the chassis backup file.
The available options are:
 - CIFS
 - NFS
3. Enter the **Network Share Address** and **Network Share Filepath**.
4. If the **Share Type** is CIFS, enter the **Domain**, **User Name**, and **Password**. Else, go to step 5.
5. In **Backup File Password**, enter the **Encryption Password** and **Confirm Encryption Password**.
The backup file is encrypted and cannot be edited.
6. In **Optional Devices**, select the compute sleds in the chassis that you want backup.
The number of selected devices is displayed in the bottom left corner of the **Backup Chassis** window.
7. Click **Backup**.
A message is displayed indicating that the backup is successful and the chassis **Overview** page is displayed.
You can check the status and details of the backup process on the **Monitoring > Jobs** page.

Restoring chassis

You can restore the configuration of a chassis using a back up file, if the backed-up configuration is of the same chassis. You must have the chassis administrator role with device configuration privilege to restore the chassis.

To restore a chassis:

1. On the chassis **Overview** page, click **More Actions > Restore**.

The **Restore Chassis** window is displayed.

2. Under **Restore File Location**, select the **Share Type** where the configuration backup file is located.
3. Enter the **Network Share Address**, and **Network Share Filepath** where the backup file is stored.
4. Enter the name of the **Backup File**.
5. If the **Share Type** is CIFS, enter the **Domain**, **Username**, and **Password** to access shared location. Else, go to step 6.
6. In the **Restore File Password** section, enter the **Encryption Password** to open the encrypted back up file.
7. Click **Restore** to restore the chassis.
A message is displayed indicating that the chassis is successfully restored.

You can check the status and details of the restore process on the **Monitoring > Jobs** page.

Exporting chassis profiles

You can export chassis profiles for cloning the settings to other chassis.

To export the chassis profile:

1. On the OME–Modular home page, click **More Actions > Export Profile**.
The **Export Profile** window is displayed.
2. Select the **Share Type**.
3. Enter the network share address and path.
4. If the **Share Type** is CIFS, enter the **Domain**, **User Name**, and **Password** to access the shared location.
5. Click **Export**.

Managing chassis failover

Failover is applicable in dual management module configuration and is the process of transferring the active role to the standby management module. Reboot the active management module and re-initialize the stand-by management module to assume the active role. The failover operation takes up to 10 minutes for completion. OME–Modular is unavailable during this process. You must have the chassis administrator privilege to start a failover.

i **NOTE:** After a failover, the chassis management performance returns to normal in a few minutes.

i **NOTE:** During a failover, the chassis power state on the OME–Modular GUI is displayed as "off". The original power state is displayed after the inventory is refreshed.

To start a failover:

On the chassis **Overview** page, click **More Actions > Failover**.

A message is displayed stating that the system cannot be accessed during a failover.

Troubleshooting in chassis

The Troubleshoot option on the OME–Modular home page enables you to use the following options to resolve issues that occur in the chassis:

- Extract Log—Use this option to extract application logs and save them to the NFS or CIFS locations on the network.
- Diagnostic Commands—Use this option to run diagnostic commands and parameters to troubleshoot the chassis network.
- Reset Management Module—Use this option to reboot the management module (MM) in a single management module configuration, and perform a failover in a dual MM configuration.
i **NOTE:** During a factory reset process, the synchronization takes about 3-5 minutes. During this period, the serial, KVM, and Quick Sync interfaces do not accept the factory password and the login attempt fails.
- Terminate Serial Connection—Use this option to end the existing serial sessions.

Blinking LEDs

You can use the **Blink LED** option on the OME–Modular home page to turn off or turn on the chassis LED.

Interfaces to access OME-Modular

After configuring the network settings in OME-Modular, you can remotely access OME-Modular using various interfaces. The following table lists the interfaces that you can use to remotely access OME-Modular.

Table 4. Management module Interfaces

Interface	Description
Web interface	<p>Provides remote access to OME-Modular using a graphical user interface. The web interface is built into the OME-Modular firmware and is accessed through the NIC interface from a supported web browser on the management station. The number of user sessions that are allowed for each interface is:</p> <ul style="list-style-type: none"> • Web interface – 6 • RESTful API – 32 • SSH – 4 <p>For a list of supported web browsers, see the Supported browsers section in the <i>OME - Modular for PowerEdge MX7000 Chassis Release Notes</i> available at www.dell.com/openmanagemanuals.</p>
Remote RACADM command-line interface	<p>Use this command-line utility to manage OME-Modular and its components. You can use remote or firmware RACADM:</p> <ul style="list-style-type: none"> • Remote RACADM is a client utility that runs on a management station. It uses the out-of-band network interface to run RACADM commands on the managed system and uses the HTTPs channel. The <code>-r</code> option runs the RACADM command over a network. • Firmware RACADM is accessible by logging in to OME-Modular using SSH or telnet. You can run the firmware RACADM commands without specifying the OME-Modular IP, user name, or password. After you enter the RACADM prompt, you can directly run the commands without the <code>racadm</code> prefix. <p>NOTE: A log for remote racadm session (login or logout) is displayed in the Audit Logs page, irrespective of the remote racadm status. However, the feature does not work if the remote racadm option is disabled.</p>
LCD	<p>Use the LCD on the front panel to:</p> <ul style="list-style-type: none"> • View alerts, OME-Modular IP or MAC address. • Set DHCP • Configure OME-Modular static IP settings. • View OME-Modular MAC address for the active MM. • View the OME-Modular VLAN ID appended to the end of MM IP, if the VLAN is already configured. • At-the-box management—Create a group, join group, leave group, or delete group. • At-the-box storage mapping resolution on compute sled replacement condition. <p>For more information about the LCD touch panel, see the <i>Dell EMC PowerEdge MX7000 Enclosure Installation and Service Manual</i>.</p>
SSH	Use SSH to connect to the MX7000 chassis and run RACADM commands locally.
RESTful API and Redfish	<p>The Redfish Scalable Platforms Management API is a standard that the Distributed Management Task Force (DMTF) has defined. Redfish is a next-generation systems management interface standard, which enables scalable, secure, and open server management. It is a new interface that uses RESTful interface semantics to access data that is defined in model format to perform out-of-band systems management. It is suitable for a wide range of servers ranging from stand-alone servers to rack mount and bladed environments and for large-scale cloud environments.</p> <p>Redfish provides the following benefits over existing server management methods:</p> <ul style="list-style-type: none"> • Increased simplicity and usability • High data security • Programmable interface that can be easily scripted

Table 4. Management module Interfaces(continued)

Interface	Description
	<ul style="list-style-type: none"> Follows widely used standards For more information, see the <i>OME and OME - Modular RESTful API Guide</i> available at www.dell.com/openmanagemanuals .
SNMP	Use SNMP to: <ol style="list-style-type: none"> Download the OME-Modular MIB file from the https://www.dell.com/support. Use MIB walker tool to get supported information using OIDs. <p>NOTE: SNMP SET is not supported.</p>
Serial	You can use the serial interface to access OME-Modular by connecting the micro USB port on the rear of the management module to a laptop and opening a terminal emulator. The user interface that is displayed enables you to log in to the management module, networking IOMs, or servers (iDRAC). You can have a maximum of one serial session open at a time.
Quick Sync	You can have a maximum of one Quick Sync session open at a time.
KVM	You can have a maximum of one KVM session open at a time.
Chassis Direct	The Chassis Direct feature enables you to access management consoles such as iDRAC and management module of devices on the MX7000 chassis.

Viewing chassis hardware

On the OME-Modular home page, click **Hardware** to view details of the hardware components that are installed in the chassis. You can also view the chassis hardware details by clicking **Devices > Chassis > View Details > Hardware**. The hardware components comprise of chassis power supplies, chassis slots, management module, fans, temperature, FRU, device management information, installed software, and management ports.

NOTE: If the Power Supply Unit (PSU) is absent, the health state and power status for the PSU are not displayed on the Chassis > Hardware > Chassis Power Supplies page.

NOTE: Maintain a minimum interval of two minutes while removing and inserting any device.

Viewing chassis alerts

On the OME-Modular home page, click **Alerts** to view details of the alerts triggered for the events that occurred in the chassis. You can also view the chassis hardware details by clicking **Devices > Chassis > View Details > Alerts**.

You can sort the list of alerts based on the following advanced filters:

- Severity
- Acknowledge
- Start Date
- End Date
- Source Name
- Category
- Subcategory
- Message

Select an alert to view the summary of the alert.

You can also perform the following activities on the **Alerts** page.

- Acknowledge**
- Unacknowledge**
- Ignore**
- Export**
- Delete**

Viewing chassis hardware logs

The logs of activities performed on the hardware components associated with the chassis are displayed on the OME–Modular **Hardware Logs** page. The log details that are displayed include severity, message ID, category, timestamp, and description. You can also view the chassis hardware logs by clicking **Devices > Chassis > View Details > Hardware Logs**.

You can perform the following tasks on the **Hardware Logs** page:

- Click **Advanced Filter** to filter the logs based on severity, message ID, start date, end date, or category.
- Click **Export > Export Current Page** to export all the displayed logs.
- Select a specific log and click **Export**.

NOTE: If a `racrestcfg` is performed, the message, "CMC8709 and CMC8710 logs are appearing 2 times each, one for slot 1 and other for slot 2", is displayed on the Hardware Logs page.

Configuring OME–Modular

The **Application Settings** menu on the home page enables you to configure various settings for OME–Modular. The settings include the following:

- Network
- Users
- Security
- Alerts

Viewing current configuration

Click **Application Settings > Network > Current Settings**.

The current network, IPv4, and IPv6 settings are displayed.

Configuring OME–Modular IP address

1. Click **Application Settings > Network > Address Configuration**.
2. Ensure that the **Enable NIC** option is selected.
3. Enable the required IP version-IPv4 or IPv6.

NOTE: The IOM and OME–Modular must be registered in the DNS. Else, the message, "Warning: Unit file of `rsyslog.service` changed on disk, 'systemctl daemon-reload' recommended.", is displayed.

NOTE: After rebooting OME–Modular, the public interface with the OME–Modular IP is available after 12 minutes approximately.

4. Enable the DHCP option, and enter the IP address and other details.

Configuring OME–Modular web server

1. Click **Application Settings > Network > Web Server Configuration**.
2. Ensure that the **Enable Web Server** option is selected.
3. Enter the timeout value in minutes.
4. Enter the port number for the web server.

You can enter a port number in the 10-65535 range. The default port number is 443.

when the web server https port settings from the lead chassis are applied to member chassis as part of the add or join member task, refresh the inventory for the lead chassis manually to see the correct https port for the member chassis, on the **Hardware > Device Management Info** page. Launch the member chassis from the lead chassis to see the port number.

If you customize the https port, OME-Modular tries to redirect to the new port automatically. However, the redirection may not work owing to security limitations of the browser. In such cases, open a new window or tab of the browser and enter the OME-Modular URL using the customized port. For example, `https://10.0.0.1:1443`

NOTE: Disabling the OME-Modular web server does not affect the launching of OME-Modular GUI on the phonebook page while using Chassis USB Direct.

NOTE: To update the webservice timeout and session configuration timeout, use the same chassis profile. Using the same chassis profile ensures that the webservice timeout and the session configuration timeout are synchronized. Else, when the webservice timeout is updated and the session configuration is processed, the session configuration overwrites the web service settings.

Configuring session inactivity timeout

1. In the **Universal Timeout** section, select the **Enable** check box and enter the time in minutes after which all the sessions must end. The duration can be in 1-1440 minutes.
If you enter the universal inactivity timeout duration, the inactivity options for the API, web interface, SSH, and Serial sessions are disabled.
2. In the **API, Web Interface, SSH, and Serial** sections, enter the time in minutes after the sessions must end and the maximum number of sessions you want to enable.
The timeout duration can be 1-1440 minutes and the maximum number of sessions can be between one and 100. The inactivity timeout duration can be 1-100 minutes for API and Serial sessions, 1-120 minutes for web interface sessions, and 1-180 minutes for SSH sessions.

The maximum number of sessions for the interfaces are as follows:

- API—1-100
- Web interface—1-6
- SSH—1-4
- Serial—1

When you downgrade from the current version of OME-Modular to an earlier version, maximum number of API sessions supported is 32. However, if you upgrade OME-Modular to the latest version that supports 100 sessions, but the API Session attribute value that is displayed is 32. You can manually set the attribute value to 100 sessions.

Configuring OME-Modular date and time settings

1. Click **Application Settings > Network > Time Configuration**.
2. Select the **Use NTP** check box, if required, and enter the NTP server details.
3. Select the required time zone.

NOTE: Any change in the attribute settings leads to IP drop or unavailability of the OME-Modular web interface for some time. However, the OME-Modular web interface recovers automatically.

Configuring OME-Modular proxy settings

1. Click **Application Settings > Network > Proxy Configuration**.
2. Select **Enable HTTP Proxy Settings**.
3. Enter the proxy address and the port number.
4. If the proxy requires authentication, select **Enable Proxy Authentication** and enter the credentials.
You can enable proxy authentication only if the **Enable HTTP Proxy Settings** option is selected.
5. Enter the proxy user credentials.

Changing device naming and preference

1. Click **Application Settings > Network > Device Name Preference**.
2. Select the naming preference.

Ports and protocols supported in OME-Modular

The table below lists the protocols and ports that are supported in OME-Modular.

Table 5. Ports and protocols that are supported in OME-Modular

Port number	Protocol	Port type	Maximum encryption level	Source	Direction	Destination	Usage
22	SSH	TCP	256-bit	External application	In	OME-Modular	Required for incoming only if FSD is used. OME-Modular administrator must enable this port only while interacting with Dell EMC.
25	SMTP	TCP	None	OME-Modular	Out	External Application	To receive email alerts from OpenManage Enterprise.
53	DNS	UDP/TCP	None	OME-Modular	Out	External Application	For DNS Queries
80	HTTP	TCP	None	External Application	In	OpenManage Enterprise Modular	The Web GUI landing page. Will redirect a user to HTTPS.
123	NTP	UDP	None	OME-Modular	Out	NTP Server	Time synchronization (if enabled).
137, 138, 139, 445	CIFS	UDP/TCP	None	OME-Modular	Out	CIFS Share	To import firmware catalogs from CIFS share.
161*	SNMP	UDP	None	External Application	In	OpenManage Enterprise Modular	For SNMP queries.
162	SNMP	UDP	None	External Application	In/Out	OpenManage Enterprise Modular	Send SNMP traps and receive Informed Request.
443	HTTPS	TCP	128-bit SSL	External Application	In/Out	OpenManage Enterprise Modular	Web GUI. To download updates and warranty information from dell.com. The 256-bit encryption is enabled while communicating with OME-Modular by using the HTTPS protocol for the web interface.

Table 5. Ports and protocols that are supported in OME-Modular(continued)

Port number	Protocol	Port type	Maximum encryption level	Source	Direction	Destination	Usage
514**	Syslog	TCP	None	OME-Modular	Out	Syslog Server	To send alert and audit log information to Syslog server
546	DHCP	TCP	None	OME-Modular	Out		Network configuration
636	LDAPS	TCP	None	OME-Modular	Out	External Application	AD/ LDAP login for Global Catalog.
3269	LDAPS	TCP	None	OME-Modular	Out	External Application	AD/ LDAP login for Global Catalog.

Legend:

- *—You can configure up to 65535 ports excluding the port number that are already allocated.
- **—Configurable ports

Configuring users and user settings

In OME-Modular, you can create up to 64 local users and assign them specific roles and privileges. Using the options available under **Application Settings > Users**, you can add and edit users, import a directory group, and view and terminate active user sessions.

 **NOTE:** You can create, delete, enable, or disable users only if you have the security setup privilege.

Viewing and editing user accounts


1. Click **Application Settings > Users**
On this page, you can view a list of users accounts and their roles, the user types, and whether the account is enabled or not.
2. Select a user and click **Edit** on the right side of the page.
3. Edit the required settings.


 **NOTE:** You can change only the password of the default "root" account.

Adding users

1. Click **Application Settings > Users**
2. Click **Add**.
3. Enter the **Username**.

The default username is "root", and you cannot edit it. You cannot disable the default account or edit the role associated with the default account. The length of the username can be 1-16 characters long and contain white spaces and alphanumeric characters. The special characters - \$, ", /, :, @, and ` are not supported.

 **NOTE:** For the OME-Modular serial interface, ensure that the length of the local or remote username does not exceed 35 characters.

 **NOTE:** Do not use "system" as a username.

4. Enter the **Password** and **Confirm Password**.

The password can be 8-32 characters long and contain at least one of the following:

- Number
- Special character—The supported special characters are - +, &, ?, >, -, }, |, ,, !, (, ' ,, _ , [, ", @, #,), *, :, \$,], /, %, =, <, :, {, |
- Uppercase letter
- Lowercase letter

5. Select a role.
6. Select **Enabled** to enable the account immediately after you create it.

 **NOTE:** For more information about the fields, see the integrated help in the OME-Modular web interface.

Enabling, disabling, and deleting users

1. Click **Application Settings > Users**.
A list of user account is displayed.
2. Select the account, and then click the required option above the list of accounts.

Recovering passwords

You must have physical access to the chassis to reset the login credentials to defaults.

Recovering passwords in single OME-Modular controller

1. From the chassis, remove the single OME-Modular controller.
2. Locate the Jumper, see the board location—P57 RESET PASSWORD, and then insert the Jumper.
3. Reinsert the controller into the chassis.
4. When OME-Modular is available, login with the username as "root" and password as "calvin".
5. After the root user authentication, change the password for the root user on the **Application Settings > Users** page.
6. Log out and log in again using the modified password to ensure that the login is successful.
7. Remove the jumper and reinsert it into the default positions—2 and 3.

Recovering passwords in dual OME-Modular controllers

1. From the chassis, remove both the OME-Modular controllers.
2. On one of the modules, locate the Jumper, see the board location—P57 RESET PASSWORD, and then insert the Jumper.
3. Reinsert only the controller, where the Jumper is installed, into the chassis.
4. When OME-Modular is available, login with the username as "root" and password as "calvin".
5. After the root user authentication, modify the password for the root user on the **Application Settings > Users** page.
6. Remove the controller where the Jumper is inserted, locate the Jumper.
7. Set the Jumper to the default position and insert the controller back into the chassis.
8. When OME-Modular is available, login with the modified password.
9. Insert the second controller to restore the MM redundancy.

User roles and privileges

Table 6. User roles and privileges

User Role	Chassis Administrator	Compute Manager	Storage Manager	Fabric Manager	Viewer
Privilege					
Viewing application information	Yes	Yes	Yes	Yes	Yes
Setting up applications such as network, NTP, and proxy	Yes	No	No	No	No
Setting up users, security login policies, and certificates	Yes	No	No	No	No
Monitoring alert policies and alert destinations	Yes	No	No	No	No

Table 6. User roles and privileges(continued)

User Role	Chassis Administrator	Compute Manager	Storage Manager	Fabric Manager	Viewer
Device power control	Yes	Yes	Yes	Yes	No
Device configuration actions Example—Applying templates, migrating profiles, and managing storage mappings	Yes	Yes	Yes	Yes	No
Updating device firmware	Yes	Yes	Yes	Yes	No
Creating and managing device templates, identity pools, and logical networks	Yes	Yes	Yes	Yes	No
Managing firmware catalogs and baseline policies	Yes	Yes	Yes	Yes	No
Power budget configuration and management	Yes	No	No	No	No

Managing user sessions

You can view and terminate existing user sessions using the **User Sessions** page, if you have the chassis administrator privilege.

Viewing user sessions

On the **Users** page, click **User Sessions**.

You can view the list and details of the users who are logged in.

Terminating user sessions

1. On the **Users** page, click **User Sessions**.
You can view the details of the users who are logged in.
2. Select the user from the list and click **Terminate**.
A message is displayed prompting you to confirm the termination.

Importing Directory Group

You can import Active Directory groups and map them to the existing OME–Modular groups.

To import the Active Directory groups:

1. On the **Users** list page, click **Import Directory Group**.
The **Import Directory** window is displayed.
2. From the **Directory Source** drop-down, select the source from which you want to import the Active Directory.
3. Under **Available Groups**, you can search for directory groups.
The list of groups is displayed below.
4. Select a group and click ">>".
The selected group is displayed under **Groups to be Imported**.
5. Click the check box corresponding to the group.
6. From the **Assign Group Role** drop-down, select the role that you want to assign to the group and click **Assign**.

Adding directory services

You can create directory services with details.

1. From the main menu, click **Application Settings > Users > Directory Services > Add**. The **Connect to Directory Service** window is displayed.

2. Select the directory type from the **Type of Directory** drop-down list.

The available options are:

- **AD**
- **LDAP**

3. Enter a name for the directory service in the **Directory Name** field.

i **NOTE:** The directory name can have a maximum of 255 characters.

4. From the **Domain Controller Lookup**, select **DNS** or **Manual**.

5. Enter the DNS domain name in the **Method** field.

i **NOTE:** If the domain controller lookup type is **Manual**, enter the Fully Qualified Domain Name (FQDN) or IP addresses of the domain controller.

- a. If you have selected the directory type as **AD**, enter the domain name in the **Group Domain** field.

i **NOTE:** This option is displayed only if the directory type is **AD**.

i **NOTE:** If the directory type is **AD**, the supported server port number is 3269 for the global catalog and 636 for domain controller. If you configure other ports for the Active Directory service, the Directory Service may not work properly as the communication with the AD server fails with different ports.

i **NOTE:** If the Server Port is 3269, the Group Domain input method is `example.com` or `ou=org, dc=example, dc=com`. And, if the Server Port is 636 or a port other than 3269, the Group Domain input method is `ou=org, dc=example, dc=com`.

- b. If you have selected the directory type as **LDAP**, enter **Bind DN** and **Bind Password** in the respective fields.

i **NOTE:** These options are displayed only if the directory type is **LDAP**.

6. Click the **Advance Options** and enter the details.

- a. If you have selected the directory type as **AD**, enter the following details:

- **Server Port** number—The server port number can be between 1 and 65535
- **Network Timeout** and **Search Timeout** in seconds
- Select the **Certificate Validation** checkbox
- Click **Select a file** to browse and upload a certificate

- b. If you have selected the directory type as **LDAP**, enter the following details:

- **Server Port** number—The server port number can be between 1 and 65535
- **Base Distinguished Name to Search**
- **Attribute of User Login**, **Attribute of Group Membership**, and **Search Filter**
- **Network Timeout** and **Search Timeout** in seconds
- Select the **Certificate Validation** checkbox
- Click **Select a file** to browse and upload a certificate

i **NOTE:** If the **Certificate Validation** check box is selected, enter the FQDN of the domain controller in the **Method** field. The certificate validation is successful only if the details of the Issuing Authority in the certificate and the FQDN match.

Deleting directory services

To delete directory services:

1. From the main menu, click **Application Settings > Users > Directory Services**.

2. Select the directory service that you want to delete and click **Delete**.

Configuring login security settings

OME–Modular supports IP range-based access restriction. You can restrict access to only a specified range of IP addresses. You can also configure lockout policies that enforce delays after certain number of failed login attempts.

Configuring login IP range

1. Click **Application Settings > Security > Login IP Range**.
2. Select **Enable IP Range**.
3. Enter the IP range in the CIDR format.

For IPv4, enter the IP address in the format—192.168.100.14/24. For IPv6, enter the IP address in the format—2001:db8::/24.

Configuring login lockout policy

1. Click **Application Settings > Security > Login Lockout Policy**.
2. Select **By User Name** to enable user account-based lockout. Select **By IP Address** to enable IP address-based lockout.
3. Enter the lockout details:
 - a. Lockout Fail Count: The number of failed login attempts. Valid values are between 2 and 16.
 - b. Lockout Fail Window: The time within which subsequent failed logins are registered. Valid time is between 2 seconds and 65,535 seconds.
 - c. Lockout Penalty Time: Time for which the logins are restricted. Valid time is between 2 seconds and 65,535 seconds.

If the IP is still unavailable, ensure that:

- The network cable is connected.
- If DHCP is configured, ensure that the cable is connected to a ToR switch that has connectivity to the DHCP server.

Enabling FIPS mode

The United States government agencies and contractors use the FIPS standards. FIPS Mode is intended to meet the requirements of FIPS 140-2 level 1.

To enable FIPS mode, click **Application Settings > Security > Federal Information Processing Standards (FIPS)**

 **NOTE:** After enabling the FIPS mode or reset configuration operation, wait for sometime for the application to become stable.

Managing certificates

You can view details of the SSL certificates on the **Certificates** page. The information includes the details of:

- The organization the certificate is issued to
- The issuing authority of the certificate
- The validity of the certificate

If you have the security setup privilege, you can perform the following tasks:

- View the SSL certificate that is deployed.
- Generate a new certificate signing request (CSR)
- Upload the server certificate, based on the CSR generated, to replace the default or currently deployed certificate.

Uploading certificates

To upload the certificate:

1. Click **Application Settings > Security > Certificates**.
2. Click **Upload** to browse and upload the certificate.

Generating certificate signing request

1. Click **Application Settings > Security > Certificates**.
2. At the bottom-right of the page, click **Generate Certificate Signing Request**.
3. Enter the required details and click **Generate**.
 - OME–Modular does not create an SSL certificate on time change or on every boot or time change and boot simultaneously.

- OME-Modular generates a new SSL certificate with validity from build_time till (build_time +10 years) only during first boot scenarios such as firmware update, racresetcfg, and FIPS mode changes.

NOTE: Only the users with the chassis administrator privileges can generate certificate signing requests.

Configuring alerts

This section allows you to configure the email, SNMP, and the syslog settings to trigger alerts.

Configuring email alerts

1. Click **Application Settings > Alerts**.
2. Click **Email Configuration**
3. Enter the **SMTP Server Network Address**.

NOTE: The SMTP server network address can have a maximum length of 255 characters.

4. If the server requires authentication, select **Enable Authentication**.

NOTE: If **Enable Authentication** is selected, you must provide the username and password to access the SMTP server.

5. Enter **SMTP Port Number**.
6. If the SMTP server is configured to use SSL, select the **SSL** option.

Configuring SNMP alerts

The SNMP alerts contain the service tag of the chassis as one of the parameters in the trap. Third-party consoles can use this information to correlate the traps with the system.

For network IOMs and compute sleds, OME-Modular subscribes to alerts through internal private VLANs—SNMP or REST. For MXG610s fiber channel switching modules, only SNMP V1 is supported and you can configure only four SNMP alert destinations.

You can configure the SNMP alert destination for IOMs from the **Application Settings > Alerts > SNMP Configuration** page. After configuring the SNMP destination, go to **I/O Settings > Replicate Alert Destinations**.

To configure SNMP alerts, perform the following steps:

1. From the main menu, select **Application Settings > Alerts**.
2. Click **SNMP Configuration**.
3. Select **Enable** to enable the configuration.
4. Enter the **Destination Address**.

You can configure up to four SNMP destinations.

5. Select the **SNMP Version**.

The available SNMP versions are:

- SNMP V1
- SNMP V2

NOTE: For MX9116n or MX5108n IOMs, only SNMP V2, is supported.

NOTE: The MX7000 chassis facilitates configuration of four SNMP destinations. However, the MXG610s FC IOM switches support only three SNMP destinations. If the fourth SNMP destination is configured, the IOM ignores it.

6. Enter the **Community String**.

When you configure the community string for SNMP V1, by default, the community string is appended with |common|FibreChannel111.

7. Select the **Port Number** and click **Send** to test the SNMP trap.

Configuring sys log alerts

You can configure up to four sys log destinations.

To configure system log alerts, perform the following steps:

1. Click **Application Settings > Alerts > Syslog Configuration**.
2. Select the **Enabled** check box corresponding to the required server.
3. Enter the destination address or the hostname.
4. Enter the port number.

Managing compute sleds

OME–Modular enables you to allocate and manage compute sleds to balance workload demands.

You can view the list and details of compute sleds on the **Compute** page. The details are—health, power state, name, IP address, service tag, and model of the chassis. You can also select a compute sled to view the graphical representation and summary of the compute sled, on the right side of the **Compute** page.

Select a compute sled from the list to view a summary of the sled on the right side. The summary includes links to launch the iDRAC and virtual consoles, name of the compute sled, device type, service tag, management IP, model, and health.

If you have the Compute Manager privileges, you can perform the following tasks in this tab:

- **Power Control** tasks:
 - **Power Off (Non-graceful)**
 - **Power Cycle System (Cold Boot)**
 - **System Reset (Warm Boot)**
 - **Power Off (Graceful)**
 - **System Reset**
 - **Power On**
- Turn-on or turn off LEDs using **Blink LED**.
- Refresh Inventory.

NOTE: When a compute sled is inserted into a chassis, sometimes the message, "No device image found", is displayed. To resolve the issue refresh the inventory of the compute sled, manually.

After performing a power operation on compute sleds, some sleds do not transition to the intended state immediately. In such cases, actual state of the compute sled is updated during the next health or inventory refresh.

NOTE: If the compute sled and fabric IOM mismatch, the health status of the compute or IOM is displayed as "Warning" in the chassis subsystem health. However, the health status is not displayed in the chassis graphical representation on the Chassis page, I/O Modules, and Compute pages.

NOTE: Occasionally, you may see messages stating the device is offline. The messages are logged when the status poll for the device indicates that the device transitioned to "off" from "on".

Topics:

- [Viewing compute overview](#)
- [Configuring compute settings](#)
- [Replacing compute sleds](#)
- [Viewing compute hardware](#)
- [Viewing compute firmware](#)
- [Viewing compute hardware logs](#)
- [Viewing compute alerts](#)

Viewing compute overview

On the compute **Overview** page, you can view a graphical representation of the compute on the left side. The compute information is displayed below the graphical representation. The information includes details such as iDRAC DNS name, model, service tag, asset service tag, express service code, management IP, system up time, populated DIMM slots, and total number of DIMM slots in the compute. You can also see the operating system and location information details.

You can also see information under the following sections:

- **Operating System Information**—Displays the name, version, and hostname of the operating system that is installed on the compute sled.
- **Location Information**—Displays the location details of the compute sled.

- **Chassis Information**—Displays the details of the chassis on which the compute sled is placed. Click **View All** to view the list of all activities in the **Jobs** page.
- **Recent Alerts**—Displays the number and details of the tasks that are performed in the compute sled. Click **View All** to view the list of all alerts that are related to the compute sled on the **Compute > Alerts** page.
- **Recent Activity**—Displays the status of the jobs that are performed in the compute sled.
- **Remote Console**—Displays a graphical representation of the remote console is displayed on the right-side of the page. Below the remote console image, you can use the following links:
 - **Launch iDRAC**—Displays the iDRAC GUI.
 - **Launch Virtual Console**—Opens the virtual console.

NOTE: The virtual console preview is unavailable for users with the "Viewer" User Role type.

- **Server Subsystems**—Displays a summary of information about the server sub systems. The information includes the health status of the components such as battery, memory, processor, and voltage.
- **Environment**—Displays the temperature and power supply information of the compute. You can also view the power and temperature statistics for the compute.

NOTE: The time that is displayed is based on the time zone of the system from where OME–Modular is accessed.

NOTE: The Launch iDRAC or Launch Virtual Console options are disabled based on the:

- **Readiness of iDRAC**
- **Power off state of the compute sled**
- **Availability of express license in iDRAC**
- **Status of firmware update in iDRAC**
- **Status of the virtual console**

Also, Internet Explorer and Safari have certain limitations that restrict the reuse of OME–Modular sessions. Hence, you are prompted to enter the OME–Modular user credentials to access iDRAC.

NOTE: The Peak Power value that is displayed is the last peak value irrespective of the power state of the device or component.

If you have the Compute Manager privileges, you can perform the following tasks in this tab:

- **Power Control** tasks:
 - **Power Off (Non-graceful)**—Turns off the server power, which is equivalent to pressing the power button when the server is turned on. This option is disabled if the server is already turned off. It does not notify the server operating system.
 - **Power Cycle System (Cold Boot)**—Turns off and then restarts the server (cold boot). This option is disabled if the server is already turned off.
 - **System Reset (Warm Boot)**—Restarts (resets) the server without turning off (warm boot).
 - **Power Off (Graceful)**—Notifies the server operating system to turn off the server. This option is disabled if the server is already turned off.
 - **System Reseat**—Removes the compute sled virtually.
 - **Power On**—Turns on the server power, which is equivalent to pressing the power button when the server is turned off. This option is disabled if the server is already turned on.
- Extract **SupportAssist** logs and reset iDRAC using **Troubleshoot**.
SupportAssist is used to collate hardware, operating system, and RAID controller logs and store the logs in the NFS or CIFS share location.
iDRAC reset helps in troubleshooting when iDRAC is noncommunicative.
- Turn-on or turn off LEDs using **Blink LED**. The available options are:
 - **1 Minute**
 - **10 Minutes**
 - **30 Minutes**
 - **1 Hour**
 - **Indefinitely**
- **Configuration Profile** tasks:
 - Associate server profiles—You can associate profiles to blade servers. The profile is extracted from the server and is attached to the slot containing the server.

- Migrate server profiles—You can migrate a profile from one server to another. The system unassigns the identity from the first server before the migration. If the unassignment fails, the system displays a critical error. You can override the error and force the migration to a new server.
- Edit server profiles—You can edit the profile characteristics that are unique to the device or slot. If a profile is attached to a compute, the updated profile configuration is propagated to the compute.
- Remove slot associations—You can detach the server profile from the slot.
- Detach server profiles—You can remove profiles that are associated with blade servers. After the server profile is detached, the identity pools are reclaimed from the MAC address pools. Detaching a profile reclaims the Identities from the device, based on the last deployed template or profile. If the last deployed template does not have the MAC Identities association, the MAC Identities that are already deployed are not reclaimed.

In MCM environment, if the compute sled on the member chassis is not reachable you can detach the profile from the lead chassis using the **Detach profile** option. The status of the **On Detach profile Reclaim Identities** task on the **Jobs** page in the lead chassis is displayed as **Completed**. However, the **On Detach profile Reclaim Identities** job fails in the member chassis.

If the compute sled on the stand-alone chassis is not reachable and you try detaching the profile, the **On Detach profile Reclaim Identities** job fails.

NOTE: The Reclaim Identities feature in OME-Modular works for both scenarios.

NOTE: When a compute sled is inserted into a chassis, sometimes the message, "No device image found", is displayed. To resolve the issue refresh the inventory of the compute sled, manually.

Configuring compute settings

You can configure the following compute settings:

- Network
- Management

Configuring compute network settings

Once Quick Deploy settings are applied to a compute sled, the settings may be reported after some time due to data refresh in OME-Modular.

To configure the compute network settings:

1. Click **Devices > Compute > View Details > Settings > Network**.
2. In the **General Settings** section, select the LAN Enablement check box to configure the network settings.
3. Configure the IPv4, IPv6, and management VLAN settings.

Configuring compute management settings

To configure the compute management settings:

1. Click **Devices > Compute > View Details > Settings > Management**.
2. Configure the password to access the iDRAC console and select **IPMI over LAN** to enable access from OME-Modular to iDRAC, through BIOS.

Replacing compute sleds

The rip-and-replace feature of OME-Modular enables you to replace a failed compute sled, storage sled, or IOM, and apply the configuration automatically.

NOTE: While replacing compute sleds, ensure that the:

- **Compute sled is turned off and the compute nodes in the chassis contain PERC or HBA controllers.**
- **SAS IOMs and storage sleds are installed in the chassis.**
- When you replace a compute sled, with a service tag, with a compute sled of another service tag, and the storage sleds are mapped to the compute node slot, the power on the particular compute sled is turned off. An option to unblock power is displayed on the **Devices > Compute > Overview** page for the compute sled.
- When you remove a compute sled, containing an HBA 330 controller with shared mappings, and replace it with a compute sled containing a PERC controller, the sled is checked to ensure that no shared mappings exist. If shared mappings exist, a message is

displayed on the **Devices > Compute > Overview** page for the compute sled, prompting you to clear the mapping. The compute sled is turned off.

- When you remove a compute sled containing a PERC controller with mappings, and replace it with a new compute sled having an HBA 330 controller with a different service tag, a message is displayed on the **Devices > Compute > Overview** page for the compute sled, prompting you to clear or accept the mapping. However, the compute sled is turned on in this scenario.

The following flowchart and table illustrate the behavior of OME-Modular and the LCD panel on the chassis when the compute sled is replaced:

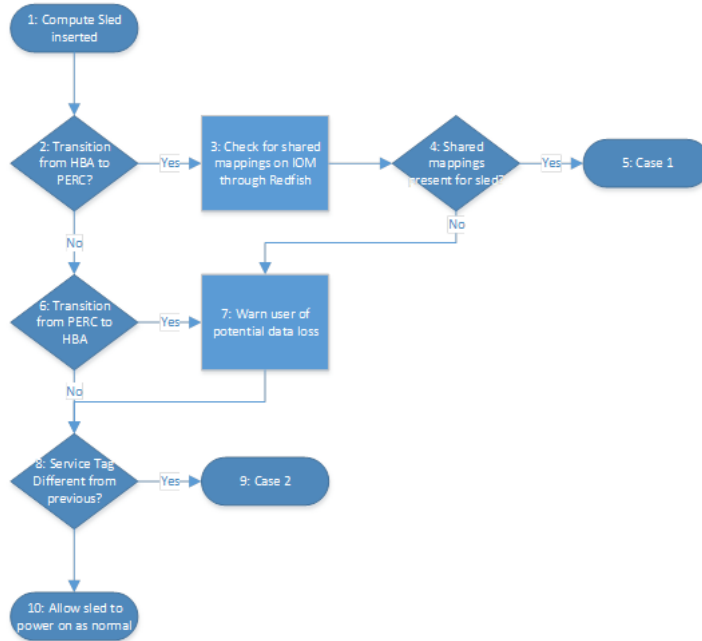


Figure 1. Compute sled replacement—flowchart

Table 7. Compute sled replacement - Behavior of OME-Modular and LCD panel

	OME-Modular behavior	LCD behavior
Case 1	Enables users to clear all mappings to the compute sled.	Enables users to clear all mappings to the compute sled.
Case 2	Enables users to clear or retain all mappings to the compute sled.	Enables users to clear or retain all mappings to the compute sled.

Viewing compute hardware

You can view the details of the hardware components that are installed in the compute sled, on the compute **Hardware** page. The hardware components include processor, storage controller, and FRU.

The deployment and configuration jobs on the compute sled are performed only for the first time, if the profile and sled device ID are unchanged. If the sled is removed and reinserted, the deployment and configuration job is not performed. This condition is applicable to the **Edit Profile** task too.

NOTE: If the storage controller cards are absent in iDRAC, the storage enclosure details are not displayed on the **Compute > View Details > Hardware > Storage Enclosure** page.

Viewing compute firmware

You can view the firmware list for the compute in the compute **Firmware** page. Click **Devices > Compute > View Details > Firmware**.

The details include name of the device or component, impact assessment, current version, and baseline version.

You can perform the following tasks on the Firmware page:

- Update the existing firmware on the compute using **Update Firmware**.

- Downgrade the updated firmware version to the previous version using **Rollback Firmware**.
- Export the firmware baseline report in a `.csv` format using **Export**.

Viewing compute hardware logs

The logs of activities performed on the hardware components associated with the compute sled are displayed on the compute **Hardware Logs** page. The log details that are displayed include severity, message ID, category, timestamp, and description.

To view the hardware logs, click **Devices > Compute > View Details > Hardware Logs**.

You can also perform the following tasks on the **Hardware Logs** page:

- Filter the logs using **Advanced Filter**—You can filter the logs based on severity, message ID, start date, end date, or category.
- Select logs and include comments for them using **Add Comment**.
- Export logs displayed on the current page or export specific logs using **Export**.

Viewing compute alerts

You can view the list of alerts and warnings for compute sleds on the **Alerts** page.

To view the compute alerts, click **Devices > Compute > View Details > Alerts**.

You can sort the list of alerts based on the following advanced filters:

- Severity
- Acknowledge
- Start Date
- End Date
- Category
- Subcategory
- Message

Select an alert to view the summary on the right side of the **Alerts**.

You can also perform the following activities on the **Alerts** page.

- **Acknowledge**
- **Unacknowledge**
- **Ignore**
- **Export**
- **Delete**

Managing storage

This chapter describes the Storage and IOM features of OME–Modular. It also provides details about performing various storage-related tasks. The SAS IOMs manage the storage enclosures. SAS IOMs facilitate communication between storage and compute sled and also help in assigning the storage to the compute sleds. You can assign storage devices as:

- Specific drive bays storage to compute sleds
- Entire storage enclosure to compute sleds

You can use the options available on the storage page to perform power operations, update firmware, manage hardware settings, and configure alerts for the storage devices.

For more information about SAS Storage, see [Managing SAS IOMs](#).

Topics:

- [Storage overview](#)
- [Viewing hardware details](#)
- [Assigning drives to a compute sled](#)
- [Assigning storage enclosure to a compute sled](#)
- [Replacing storage sleds](#)
- [Updating enclosure firmware](#)
- [Downgrading storage enclosure firmware](#)
- [Managing SAS IOMs](#)

Storage overview

On the **Storage Overview** page, you can view all the storage enclosures that are installed in the chassis. You can also perform a virtual reset of the storage enclosure and blink the LEDs to identify the storage enclosures.

To view the available storage enclosures or sleds:

1. From the **Devices** drop-down menu, select **Storage**.
2. Select the storage sled from the list of the storage devices.
3. Click **View Details**.

The storage **Overview** page is displayed.

Performing a storage system reset

You can perform a system reset remotely using the OME–Modular. The system reset option simulates a physical sled removal and reinstallation.

To perform system reset:

1. From the **Devices** drop-down menu, select **Storage**.
2. Select the storage sled you want to reset.
3. Click **Power Control** and click **System Reset**.
4. Click **Confirm**.

 **NOTE:** The storage sled, if assigned to compute sleds that are powered on, causes input/output disruption.

Blinking LED

You can locate a storage sled within a chassis by making the sled LED blink. This is helpful in identifying a system. To turn on the LED blinking:


1. From the **Devices** drop-down menu, select **Storage**.

2. Select the storage sled.
3. Click **Blink LED** and click **Turn On**.

To turn off the LED blinking:

1. From the **Devices** drop-down menu, select **Storage**.
2. Select the storage sled.
3. Click **Blink LED** and click **Turn Off**.

You can pull out the storage sled trays from the chassis to access the storage sled drives. When a tray is opened, the storage sled drive is away from the chassis and does not have cooling support causing the temperature of the drive to reach a critical level. When the tray is opened, the LCD displays count down timer starting from five minutes. Close the tray within five minutes for cooling the storage drive. Also, if another tray containing a storage sled drive is opened, the current warning display is not affected. You can dismiss the LCD warning display.

 **NOTE:** The LCD display for storage-mapping owing to server replacement takes priority over opening of the storage tray. If LCD has completed displaying storage-mapping menus and a storage tray is still open, a warning, stating the storage tray is open, is displayed.

Editing storage sled assignments

You can change the assignments of the device using **Edit Assignments** option. To edit assignments:

- On the storage **Overview** page, click **Edit Assignments**.
The **Hardware** page is displayed.
- Select the hardware component, and change the assignment. For more information, see [Assigning drives to a compute sled](#).

Other information

On the **Hardware** page, you can view more information about the device as follows:

- **Storage Enclosure Information**—Provides the information of an enclosure, such as **Name**, **FGDD**, **Model**, **Service Tag**, **Asset Tag**, **Power State**, **Firmware Version**, **Drive Slot Count**, and **Assignment Mode**
- **Chassis Information**—Provides the information of a chassis, such as **Chassis**, **Slot Name**, and **Slot**
- **Connected I/O Module Information**—Provides the information of an I/O module such as **I/O Module Name** and **Multipath**
- **Recent Alerts**—Provides the list of the recent alerts
- **Recent Activity**—Provides the list of recent activities
- **Storage Subsystems**—Provides the list of storage subsystem
- **Environment**—Provides the information of the power usage

Viewing hardware details

The hardware components of a storage sled comprise of hard drives, enclosure management modules (EMMs), Field Replaceable Unit (FRUs), and installed software. To view the details of hardware components in the storage sled:


1. From the **Devices** drop down, select **Storage**.
2. Select a storage from the list of the storage devices.
3. On the right side, click **View Details**.
4. To view the hardware details, click **Hardware**. The hardware components in the storage sled are displayed at the top of the **Hardware** page.

Viewing drive details

To view the list of drives in the storage sled, click **Hardware > Hard Drives**. You can assign individual hard drive to a compute sleds. You can update firmware of these drives using iDRAC web interface.

Current Mode—Indicates if the hard drive is assigned to an enclosure or to a single compute node slot.

- **Enclosure-Assigned**—In this mode, you can assign an entire storage sled to one or more compute node slot.

 **NOTE:** You cannot assign storage when a redundant SAS IOM setup is temporarily degraded to non-redundant state.

NOTE: The storage enclosure is assigned to the slots of the compute slots and not to the sled itself. If a compute sled is replaced with another sled on the same slot, then the storage enclosure gets assigned to the new sled automatically. However, if you move the compute sled from one slot to another, you must remap the storage to that sled.

Drive-Assigned—In this mode, you can select a hard drive slot and assign it to a compute node slot.

CAUTION: Assigning a hard drive to a compute node slot may result in loss of data.

Assigning drives to a compute sled

Using the **Drive-Assigned** mode, you can map the drives in a storage enclosure to a compute sled slot. If the compute sled fails, the drive remains assigned to the slot. If the sled is moved to another slot on the chassis, reassign the drives to the new slot. To configure RAID on the drives, use the iDRAC web interface, a server configuration profile, or an operating system deployment script, after the drive assignment is complete.

CAUTION: Before you assign a drive to a slot, ensure that the data from the drive is backed up.

NOTE: The HBA330 controller card does not set a status for the hard drives when the hard drives are removed from the storage sleds after the hard drives are assigned to compute sleds.

To assign a drive:

1. From the **Devices** drop-down, select **Storage**.
2. Select the storage sled from the list of the storage devices.
3. Click **View Details**.
The storage **Overview** page is displayed.
4. Click **Hardware**.
The drive list is displayed.

NOTE: Ensure that the **Drive-Assigned** mode is selected.

5. Select one or more drives and click **Assign Drive to Slot**.
The **Assign Hard Drive to Compute** page is displayed.
6. Select the slot and click **Assign**.

When a drive is reassigned from one compute sled to another, the enclosure status and spin-up state of the drive is the same. If a drive is in power-savings mode, the status of the drive is displayed as "starting".

Assigning storage enclosure to a compute sled

Using the **Enclosure-Assigned** mode you can assign a storage enclosure to one or more compute sleds with HBA330 mini-mezzanine adapter. Using this mode, you can also assign a storage enclosure to an empty slot. If the sled is removed and installed to another slot the assignment must be performed again.

CAUTION: Before you assign an enclosure to a slot, ensure that the data from the drive is backed up.

NOTE: Systems with H745P MX controller only support a single storage enclosure mapping.

To assign an enclosure:

1. From the **Devices** drop-down list, select **Storage**.
2. Select the storage sled from the list of the storage devices.
3. Click **View Details**.
The storage **Overview** page is displayed.
4. Click **Hardware** and select **Enclosure-Assigned**.
A warning message about loss of data while selecting this mode is displayed.
5. Select **I understand that resetting this assignment could result in data loss** and click **Ok**.
6. Select the compute sled slots and click **Assign**.

After replacing PERC card wait for some time for OME-Modular to get the new inventory details from iDRAC before performing the assignment operation. Else, refresh the inventory on the **Compute** page, manually.

Replacing storage sleds

When you remove a storage sled from one slot and insert it into another slot on the chassis, the mapping on the new slot is used for the storage sled. If you replace the storage sled with a brand new sled which does not have a service tag, the service tag and the mapping of the sled that was present in the slot earlier, are applied. However, the storage sled firmware is not replaced automatically.

Updating enclosure firmware

You can update or rollback the storage enclosure firmware using the OME–Modular. Use the following methods to update the firmware:

1. Dell Update Package (DUP)
2. Catalog-based compliance method

 **NOTE:** The OME–Modular is inaccessible during the update process.

Updating the firmware using DUP

1. Download the DUP from www.dell.com/support/drivers.
2. On the OME–Modular web interface, go to **Devices > Storage**.
3. Select the storage sled on which you want to update the firmware.
4. Click **Update Firmware**.
5. Select the **Individual package** option and click **Browse** to go to the location where you have downloaded the DUP. Wait for the comparison report, the supported components are displayed.
6. Select the required components and click **Update** to start the firmware update.
7. Go to the **Monitoring > Jobs** page to view the job status.

Updating the firmware using catalog-based compliance

1. On the OME–Modular web interface, go to **Devices > Storage**.
2. Select the storage sled on which you want to update the firmware.
3. Click **Update Firmware**.
4. Select the baseline and click **Next**.
The Schedule Update page is displayed.
5. Select the **Schedule Update** options as required.
 - **Update Now**—apply the firmware updates immediately.
 - **Schedule Later**—schedule the firmware updates for a later date. Select the required date and time.

Downgrading storage enclosure firmware

Follow these steps to roll back the firmware for a storage enclosure:

1. On the OME–Modular web interface, go to **Devices > Storage**.
2. Select the system and click **View Details**.
3. Click **Rollback Firmware**.
4. Select the available version of the firmware and click **Confirm** to continue.

Managing SAS IOMs

The internal connection of the storage subsystem is called "Fabric C", which serves as a communication mode between compute sleds and storage enclosures. The "Fabric C" is used for SAS of FC storage connectivity and includes a midplane. SAS IOMs allow creating storage assignments in which you can map storage enclosure drives or whole storage enclosures to compute sleds. SAS IOMs provide multi-path input out access for compute sleds to drive elements. The active module manages the SAS IOM and is responsible for all inventory and storage assignments on the fabric.

A single width compute sled can support one Fab-C mezzanine card that connects to each IOM through a x4 link. Each lane in the link supports 12Gbps SAS for a total of 48Gbps link to each SAS IOM. In SAS IOMs, the Fab-C IOMs are used to provide SAS switching between compute sleds and internal storage sleds such as PowerEdge MX5016s.

For information on the tasks that you can perform on the I/O Modules page for SAS, see [Managing IOMs](#).

SAS IOM Overview

The SAS IOM **Overview** page displays details of the IOM, chassis, the list of recent alerts, and recent activities. The IOM information consists of the model, power state, firmware version, fabric type, and management role of the IOM. The management roles can be of three types:

- Active
- Passive
- Degraded

A healthy system has one "Active" and one "Passive" SAS IOM.

The chassis information consists of the name of the chassis, slot name, and slot number.

Information about the SAS IOM storage subsystems is also displayed on the right side of the **Overview** page. The storage subsystem information consists of the name of the subsystem and health status. Click **View Details** to view the alerts and alert details. The details consist of the message ID, message, timestamp when the alert was triggered, and recommended action.

To view the IOM overview:

1. From the menu bar, click **Devices > I/O Modules**. The **I/O Modules** list page is displayed.
2. Select the IOM whose details you want to view. A summary of the selected IOM is displayed on the right side. The summary consists of the name of the IOM, device type, management IP, model, health status, and availability.
3. Click **View Details**. The **Overview** page is displayed.

On the **IOM Overview** page, you can perform the following tasks:

- Power Control—Turn on, turn off, power cycle, or system reset operations.
 - Turn on or turn off—When you turn off the IOM, the status of the IOM is "Offline". As a result, status of the peer IOM may be "Degraded". When you power cycle the IOM, it causes a warm reboot of the IOM.
 - Power Cycle—The Power Cycle option initiates a warm reboot of the IOM. In this instance, the power is not removed from the IOM and the core systems of the IOM reboot.
 - System Reset—The System Reset option initiates a cold reboot of the IOM. In this instance, the power is removed from the IOM and the IOM reboots.
- NOTE:** After the power reset of the SAS IOM, the IOM turns on within a minute. Any mismatch in the power status of the IOM is corrected through refreshing the inventory or is corrected automatically with the default inventory task.
- Blink LED—Turn on or turn off to identify the IOM LEDs.
 - Clear Configuration—Delete the storage IOM configuration.
 - Extract Log—Extract the IOM activities log to a CIFS or NFS share location.
 - View a list of the latest alerts and the date and time when the alerts were generated, in the **Recent Alerts** section. To view a list of all alerts click **View All**. The **Alerts** page with all alerts that are related to the IOM is displayed.
 - View a list of all activities that are related to the IOM, the rate of completion of the activity, and the date of time when the activity began, in the **Recent Activity** section. To view a list of all activities that are related to the IOM, click **View All**. The **Jobs** page with a list of all the jobs that are related to the IOM is displayed.
 - View the power statistics of the IOM by clicking **View Power Statistics** in the **Environment** section. The statistics comprise peak power timestamp, minimum power timestamp, date, and time from when the statistics is recorded. Click **Reset** to reset the power statistics data.

NOTE: If you perform the Clear operation on a SAS IOM, the IOM becomes active, if it is not already active and, the storage configuration on both the SAS IOMs is cleared.

NOTE: Resolve any suboptimal health of the IOM, other than firmware mismatch, before updating the firmware. This action ensures that the firmware is updated without downgrading the health of the SAS IOM.

Force active

You can use **More Actions > Force Active** to perform a failover on a "Passive" or "Degraded" switch. Performing a "Force Active" operation on the SAS IOM is considered a disruptive operation and must only be used if necessary. When you perform a "Force Active" operation, the SAS IOM becomes "Active" and the associated storage configuration is applied to the chassis.

You can use the **Force Active** option to resolve mismatches that occur when:

- The switches were configured earlier but are inserted in a chassis that did not have SAS IOMs earlier.
- Two switches from two different chassis are inserted into a third chassis.

You can also use **Force Active** as a preemptive action for servicing a switch. Ensure that the remaining switch is "Active" before removing the switch that must be serviced. This in turn, prevents any disruption to the fabric that might occur if the switch is removed when the other switch is "Passive".

Clearing configuration

You can clear the storage configuration of the SAS IOMs using **More Actions > Clear**. When you click **Clear**, the SAS IOM becomes "Active" and the storage configuration is cleared from the chassis.

You can use the **Clear** option to:

- Reset a chassis configuration in one step.
- Resolve a worst case mismatch where two switches are taken from two different chassis and inserted into a third chassis. In this scenario, it is unlikely that the two switches have a correct configuration. It is recommended that the **Clear** option is used to clear existing configuration and create a correct configuration.

Use the **Force Active** and **Clear** options to act upon some critical and warning messages that are displayed in the OME–Modular web interface, particularly, for a configuration mismatch.

Extracting IOM logs

You can collect a log bundle for support by selecting **Extract Log**. The log bundle collected from the SAS IOM also contains the associated logs from all storage enclosures that are discovered by the IOM even if they are not currently present in the chassis.

Managing templates

OME–Modular allows you to configure servers based on templates. A server template is a consolidation of configuration parameters extracted from a server and used for replicating the configuration to multiple servers quickly. A server profile is a combination of template and identity settings that are applied to a specific or multiple servers, or saved for later use.

You must have the template management privilege to create templates. A server template comprises the following categories:

- iDRAC configuration—Configuration specific to iDRAC
- BIOS configuration—Set of BIOS attributes
- Storage configuration—Internal storage configuration
- NIC configuration—Configuration of NICs

To view the list of existing templates, click **Configuration > Deploy**. The **Deploy** page is displayed.

You can sort the list of templates based on the name and status of the template.

On this page, you can perform the following tasks:

- Create templates
- Edit templates
- Clone templates
- Export templates
- Delete templates
- Edit network
- Deploy template

Topics:

- [Viewing template details](#)
- [Creating templates](#)
- [Deploying templates](#)
- [Editing templates](#)
- [Editing template networks](#)
- [Cloning templates](#)
- [Exporting templates](#)
- [Deleting templates](#)

Viewing template details

To view the template details.

1. On the **Deploy** page, select the template of which you want to view the details. A summary of the template is displayed on right side.
2. Click **View Details**. The **Template Details** page is displayed.

The details that are displayed are—name and description of the template, timestamp when the template was last updated, and the name of the user who last updated it. You can also view the configuration details such as server profile and BIOS information.

You can perform the following tasks on the **Template Details** page:

- Deploy the template
- Edit the template details

Creating templates

You can create templates in the following ways:

- Clone from an existing server—**Reference Device**

- Import from an external source—**Import from File**

To create a template from a reference device:

1. On the **Deploy** page, click **Create Template** and select **From Reference Device**.
The **Create Template** wizard is displayed.
2. Enter the name and description for the template and click **Next**.
The **Reference Device** tab is displayed.
3. Click **Select Device** to view the **Select Devices** window where you can select the device or chassis based on which you want to create the template.
4. Select the configuration elements that you want to clone.

Importing templates

To import an existing template:

1. On the **Deploy** page, click **Create Template** and select **Import from File**.
The **Import Template** window is displayed.
2. Enter a name for the template and **Select a file** to go to the location where the template that you want to import is stored.

Deploying templates

You can create server profiles from templates by entering identity information that is unique to each server. The information includes input output identity information and system-specific attributes such as NIC, RAID, iDRAC, or BIOS information. You can deploy templates from the **Deploy** and **Template Details** pages.

After a template is deployed on one or more servers along with VLAN configurations, if you make a mistake or decide to change the existing VLAN configurations on the Fabric Manager, then you must perform the deployment workflow again. In the deployment workflow, server is deployed after when the VLAN is configured on the Fabric Manager.

The system specific attributes that are defined in the template are not deployed automatically. Redefine the attributes for the target system that is selected for the deployment. Use **Quick Deploy** to set the VLAN ID for the system.

When you deploy an imported template where NPAR is enabled, it does not configure the bandwidth settings on fabric mode IOMs.

To deploy a template from the **Deploy** page:

1. Select the required template and click **Deploy Template**.
If the template has identity attributes, but is not associated with a virtual identity pool, a message is displayed that the physical identities are used for the deployment. Else, the **Deploy Template** wizard
2. Select the target device on which you want to deploy the template, configure the iDRAC management IP settings, and schedule the deployment.

Deploying templates from Template Details page

To deploy a template from the **Template Details** page:

1. On the **Template Details** page, click **Deploy Template**.
If the template has identity attributes, but is not associated with a virtual identity pool, a message is displayed that the physical identities are used for the deployment. Else, the **Deploy Template** wizard is displayed.
2. Select the target device on which you want to deploy the template, configure the iDRAC management IP settings, and schedule the deployment.

Editing templates

You can only modify the name and description of the template from the **Deploy** and **Template Details** pages.

1. On the **Deploy** page, select the template that you want to modify and click **Edit**. Else, on the **Template Details** page, click **Edit**.
The **Edit Template** window is displayed.
2. Make the required changes.

Editing template networks

To edit template network details:

1. On the **Deploy** page, select the template whose network details you want to modify and click **Edit Network**. The **Edit Network** window is displayed.
2. Modify the **Identity Pool**, if necessary.
3. Select the NIC grouping option for the port.

NIC teaming is suggested for redundancy, though it is not required. NIC Partitioning (NPAR) can impact how NIC teaming operates. Based on restrictions that are related to NIC partitioning, which NIC vendors implement, certain configurations prevent certain types of teaming. The following restrictions are applicable to both Full Switch and SmartFabric modes:

- If NPAR is not used, both Switch-dependent (LACP) and Other (Switch-independent) teaming methods are supported.
- If NPAR is used, only Other (Switch-independent) teaming methods are supported. Switch-dependent teaming is not supported.

The NIC teaming feature is applicable to IOM versions 10.5.0 and later.

Refer to the network adapter or operating system documentation for detailed NIC teaming instructions.

The available NIC grouping options are:

- No Teaming—NICs are not bonded and provide no load balancing or redundancy.
 - LACP—Also referred to as Switch Dependent, 802.3ad or Dynamic Link Aggregation. The LACP teaming method uses the LACP protocol to understand the teaming topology. It provides Active-Active teaming with load balancing and redundancy. With this option, only the native VLAN is programmed on non-LAG interfaces. All tagged VLANS wait until the LACP LAG is enabled on the NICs. The following restrictions are applicable to LACP teaming:
 - The IDRAC shared LOM feature can only be used if “Failover” option on IDRAC is enabled.
 - If the host operating system is Windows, the LACP timer must be set to “slow” (also referred to as “normal”).
 - Other—Refers to a NIC teaming method where the switch is unaware of the teaming technology that is used. The "other" option involves using the operating system and NIC device drivers on the server to team the NICs. Each NIC vendor may provide slightly different implementations with different pros and cons.
4. Select the tagged and untagged VLANs, if necessary.

Cloning templates

To create a copy of a template:

On the **Deploy** page, select the template of which you want to create a copy, and click **Clone**.

Exporting templates

You can export the templates to a network share or a local drive on your system.

To export a template:

On the **Deploy** page, select the template that you want to export and click **Export**.

A message is displayed to confirm the export action. The template is exported in `.xml` format to a local drive on your system or a network share.

Deleting templates

To delete templates:

1. On the **Deploy** page, select the template that you want to delete and click **Delete**. A message is displayed prompting you to confirm the deletion.
2. Click **Yes** to proceed.

When a template is deleted, the unassigned identity pools in the template are restored to the identity pool.

Managing identity pools

Identity pools are used in template-based deployment of servers. They facilitate virtualization of network identities that are required for accessing systems using Ethernet, iSCSI, FCoE, or Fibre Channel (FC). You can enter the information that is required for managing the I/O identities. The identities, in turn, are managed by chassis management applications such as OME-Modular.

When you start a server deployment process, the next available identity is fetched from the pool and used for provisioning a server from the template description. You can migrate the server profile from one server to another without losing access to the network or storage resources.

You can also associate server profiles with slots. The server profile uses the reserved identity from the pool to provision a server.

You must have the template management privilege to manage identity pools. An identity pool contains a name, description, and category. The category can be of the following types:

- Ethernet
- iSCSI
- FCoE
- FC

To view the list of identity pools, click **Configuration > Identity Pools**.

The **Identity Pools** page is displayed with the list of available identity pools and their key attributes. You can perform the following tasks on the **Identity Pools** page:

- View the summary and usage details of the identity pool
- Create identity pools
- Edit identity pools
- Delete identity pools
- Export identity pools

Select an identity pool to view the summary and usage details of the identity pool. You can sort the usage details by selecting the category of the identity pool.

For Intel NICs, all partitions on a port share the same IQN. Hence, a duplicate iSCSI IQN is displayed on the **Identity Pools > Usage** page when the **View By** option is iSCSI.

You can also use the RESTful API commands to create and edit identity pools.

 **NOTE: The Identity Pools page displays the MAC association even if the deployed template for the destination device is deleted.**

Topics:

- [Creating identity pools](#)
- [Editing identity pools](#)
- [Exporting identity pools](#)
- [Deleting identity pools](#)

Creating identity pools

You can create up to 4096 MAC addresses in an identity pool. An error message is displayed when:

- There are errors such as overlap in identity values with an existing pool.
- Syntax errors while entering the MAC, IQN, or network addresses.

Each identity pool provides information about the state of each identity in the pool. The states could be:

- Assigned
- Reserved

If the identity is assigned, the information about the assigned server and NIC Identifier is displayed. If the identity is reserved, the information about the assigned slot in the chassis is displayed.

You can create an identity pool with only the name and description and configure the details later.

 **NOTE:** You can clear identities by disabling the I/O Identity Optimization option in iDRAC.

To create identity pools:

1. Click **Configuration > Identity Pools**.
The **Identity Pools** page is displayed with the list of available identity pools and their key attributes.
2. Click **Create**.
The **Create Identity Pool** wizard is displayed.
3. Enter a name and description for the identity pool and click **Next**.
The **Ethernet** tab is displayed.
4. Select **Include Ethernet virtual MAC Addresses** to enter the **Starting MAC Address**, select the **Number of Virtual MAC Identities** you want, and click **Next**.

The MAC addresses can be in the following formats:

- AA:BB:CC:DD:EE:FF
- AA-BB-CC-DD-EE-FF
- AA.BB.CC.DD.EE.FF

You can choose to create the identity pools from iSCSI, FCoE, or FC.

The **iSCSI** tab is displayed.

5. Select the **Include iSCSI MAC Addresses** to enter the **Starting MAC Address** and select the **Number of iSCSI MAC addresses** or IQN addresses you want.
6. Select the **Configure iSCSI Initiator** to enter the **IQN Prefix**.
The pool of IQN addresses is generated automatically by appending the generated number to the prefix in the format—<IQN Prefix>.<number>
7. Select the **Enable iSCSI Initiator IP Pool** to enter the **IP Address Range, Gateway, Primary DNS Server, Secondary DNS Server**, and select the **Subnet Mask**.

The iSCSI Initiator IP settings are used only when the iSCSI is configured for booting, and when the iSCSI Initiator configuration through DHCP is disabled. When iSCSI Initiator configuration through DHCP is enabled, all these values are obtained from a designated DHCP server.

The IP Address Range and Subnet Mask fields are used to specify a pool of IP addresses that OME–Modular can assign to a device. The device can use the IP in the iSCSI Initiator configuration. Unlike the MAC address pools, a count is not specified for the IP Address Range. The pool of IP addresses can also be used to generate the initiator IP. OME–Modular supports the IPv4 format of IP addresses range in the following formats:

- A.B.C.D - W.X.Y.Z
- A.B.C.D-E, A.B.C.
- A.B.C.D/E—This format is a Classless Inter-Domain Routing (CIDR) notation for IPv4.

A maximum of 64,000 IP addresses is allowed for a pool.

OME–Modular uses the Gateway, Primary DNS, and Secondary DNS server values while deploying a template instead of using the values in the template. OME–Modular does not assign the Gateway, Primary DNS, and Secondary DNS server values from the IP address pool, if the values are within the specified IP address range. The Gateway, Primary DNS, and Secondary DNS server values serve as exclusions from the specified IP Address Range, when applicable.

8. You can select the **Include FCoE Identity** to enter the **Starting MAC Address** and select the number of **Number of FCoE Identities** you want.

The WWPN/WWNN values are generated from the MAC address. The WWPN address is prefixed with 0x2001 while the WWNN address is prefixed with 0x2000. This format is based on an algorithm similar to FlexAddresses.

9. Select the **Include FC Identity** to enter the **Postfix (6 octets)** and select the **Number of WWPN/WWNN Addresses**.

Editing identity pools

You can modify the number of entries in the identity pool. However, you cannot reduce the size of the identities that are already assigned or reserved. For example, in a pool of 100 MAC addresses, if 94 of the addresses are assigned or reserved, you cannot reduce the number of MAC addresses to less than 94.

To edit an identity pool:

1. On the **Identity Pools** page, select the identity pool and click **Edit**.

The **Edit Identity Pool** window is displayed.

2. Make the required changes.

Exporting identity pools

You can export the identity pools in a `.csv` format to a network share or local drive on your system.

To export identity pools:

On the **Identity Pools** page, select the identity pools and click **Export**.

Deleting identity pools

You can delete identity pools that are not assigned or reserved. When you attempt deleting identity pools that are associated with templates, a warning message is displayed.

To delete identity pools:

On the **Identity Pools** page, select the identity pools that you want to delete and click **Delete**.

Ethernet IO Modules

The MX7000 supports the following Ethernet I/O Modules (IOMs):

- Managed Ethernet switches:
 - MX9116n Fabric Switching Engine
 - MX5108n Ethernet Switch
- Unmanaged devices:
 - MX7116n Fabric Expander Module
 - PowerEdge MX 25Gb Ethernet Pass-Through Module
 - PowerEdge MX 10GBASE-T Ethernet Pass-Through Module

Ethernet IOMs are supported in Fabrics A and B. For details about the supported IOM slots, see [Supported slot configurations for IOMs](#).

The Ethernet switches operate in two modes:

- Full Switch mode (default)
- SmartFabric Services mode or Fabric mode

By default, an Ethernet switch operates in Full Switch mode.

In Full Switch mode, the switch operates as a full L2/L3 switch with all functionality supported by the OS10 and the underlying hardware. The switch configuration is done through the CLI. For information about configuring a switch using the CLI, see the *OS10 Enterprise Edition User Guide*

You can use OME–Modular to perform the following tasks:

- Configure host name, SNMP, and NTP settings.
- Configure port breakout modes.
- Set ports up or down.
- Monitor health, logs, alerts, and events.
- Update and manage firmware.
- View the physical topology.
- Perform power control operations.

It is recommended that you use the full switch mode when you require a feature or network architecture that is unavailable with SmartFabric Services.

For information on Fabric mode, see [SmartFabric Services](#).

Managing Ethernet IOMs

The **I/O Modules** page displays the health and asset information of the IOMs. If you have the fabric manager role with device configuration and power control privileges, you can perform the following tasks on the **I/O Module** page:

- Power Cycle—Turn on, turn off, or perform a system reset on the IOM
- Update firmware, if applicable
- Blink LED—Turn on or turn off the IOM Identification LED.
- Refresh Inventory

You must have the device configuration privileges to set up network IOMs and perform configuration tasks on them.

NOTE: When a switch changes between Full Switch and Fabric modes, it reboots.

NOTE: If the compute sled and fabric IOM mismatch, the health status of the compute or IOM is displayed as "Warning" in the chassis subsystem health. However, the health status is not displayed in the chassis graphical representation on the Chassis page, I/O Modules, and Compute pages.

Topics:

- [Viewing hardware details](#)
- [Configuring IOM settings](#)

Viewing hardware details

You can view information for the following IOM hardware:

- FRU
- Device Management Info
- Installed Software
- Port Information

NOTE: If the physical port is added as part of the port channel, it is listed under the port channel group instead of the physical port.

NOTE: The URL attribute is displayed as "N/A" on the Hardware > Device Management Info page for FC IOMs, owing to device capability limitations.

For **Port Information**, when you enable auto-negotiation, peer devices exchange capabilities such as speed and settle on mutually acceptable configuration. However, when the auto-negotiation is disabled, the peer devices may not exchange capabilities. Hence, Dell EMC recommends that the configuration on both peer devices is identical.

The guidelines for auto-negotiation process are as follows:

- MX9116n, MX7116n, and MX5108n IOMs support only 25G speeds on server facing ports.
- By default, auto-negotiation is enabled on server facing 25G ports, as per the IEEE 802.3 Standard mandate.
- You can enable or disable auto-negotiation, but cannot configure speed on server facing ports.
- When auto-negotiation is enabled, Ethernet switches display speed capability of only 25G.

To view the hardware details:

Click **I/O Modules > View Details > Hardware** .

Configuring IOM settings

If you have the IOM device configuration privilege, you can configure the following settings for the MX9116n FSE and MX5108n Ethernet Switch IOMs:

- Network
- Administrator password
- SNMP
- Time

You must have the network administrator privilege to configure the public management IP for the IOMs. The public IP facilitates use of the IOM Command Line Interface (CLI) to configure and troubleshoot the IOMs.

Configuring IOM network settings

The network settings for IOMs include configuring the public management IP for the selected management port.

To configure the networking settings:

1. Click **All Devices > I/O Modules > View Details > Settings > Network** or **Devices > I/O Modules > View Details > Settings > Network**.

2. In the **IPv4 Settings** section, select **Enable IPv4**.

3. Enter the **IP Address**, **Subnet Mask**, and **Gateway** for the management port.

The **IP Address**, **Subnet Mask**, and **Gateway** options are enabled only if the **Enable DHCP** check box is cleared.

NOTE: For MX5108n and MX9116n IOMs, the default prefix length of the DHCP IP is 128 bits, though the DHCP server is configured for 64 bits.

4. In the **IPv6 Settings** section, select **Enable IPv6**.

5. Enter the **IPv6 Address**, select the **Prefix Length**.

The **IPv6 Address**, **Prefix Length**, and **Gateway** options are enabled only if the **Enable Autoconfiguration** check box is cleared.

6. Enter the **Gateway** for the management port.

The **IPv6 Address**, **Prefix Length**, and **Gateway** options are enabled only if the **Enable Autoconfiguration** check box is cleared.

NOTE: For tagged or untagged VLAN network, any IPv6 setting configured using OME - Modular may not have the default gateway. To get the default gateway, go to the respective OS10 CLI and enable Stateless Address Autoconfiguration (SLAAC) on the respective tagged or untagged VLAN.

7. In the **DNS Server Settings** section, enter the **Preferred DNS Server**, **Alternate DNS Server 1**, and **Alternate DNS Server 2** addresses.

For MXG610s IOMs, you can set the Preferred DNS and Alternate Server 1 and 2 addresses. However, the server address for **Alternate DNS Server 2** is not applied though the response is successful as, MXG610s IOMs support only two server address for DNS settings.

8. In the **Management VLAN** section, select **Enable VLAN** and enter the **VLAN ID**.

For MXG610s FC IOMs, DHCP works only without VLAN while Static IP works with or without VLAN configuration. To change the IP configuration from DHCP IP to Static IP, perform the following steps:

- a. Disable DHCP, configure the static IP, and save the configuration.
- b. Enable VLAN, configure the VLAN ID, and save the configuration.

Configuring Linux administrator password

The Linux administrator password is used only for troubleshooting using the Linux operating system 10 shell.

To configure the Linux administrator password:

1. Click **All Devices > I/O Modules > View Details > Settings > Management** or **Devices > I/O Modules > View Details > Settings > Management**.
The **I/O Modules** page is displayed.
2. Enter the **Host Name** and **Root Password** for the IOM.

NOTE: The Linux administrator password is used only for troubleshooting using the Operating System 10 Linux shell.
To change the OS10 admin password, connect SSH to the switch and login as "admin" using the default password, "admin". The system prompts you to change the admin password.

Configuring SNMP settings

To configure the SNMP settings:

1. Click **All Devices > I/O Modules > View Details > Settings > Monitoring** or **Devices > I/O Modules > View Details > Settings > Monitoring**.
2. Select **Enable SNMP** to configure the SNMP version and community string.

Configuring advanced settings

To configure the advanced IOM settings:

1. Click **All Devices > I/O Modules > View Details > Settings > Advanced** or **Devices > I/O Modules > View Details > Settings > Advanced**.
2. Select the options to replicate the chassis time and alert settings to the IOM.

Configuring ports

In SmartFabric mode, you can configure breakout and admin status, and MTU size for IOMs. You can configure port breakout only for port groups.

NOTE: Ensure that the peer FC port has a fixed speed and matches the speed of the IOM FC port for the link to come up.

To configure breakout:

1. Click **Devices > I/O Modules > View Details > Hardware > Port Information**.
2. Select the port group and click **Configure Breakout**.

The **Configure Breakout** window is displayed.

3. Select the **Breakout Type**.

First apply "Hardware Default" and then select the required breakout.

 **NOTE: Breakouts can be configured only for Fabric Mode IOMs.**

Configuring admin status

You can switch the admin status for all ports, which is enabled by default. For the MX9116n FSE port groups 1/1/15 and 1/1/16, when you breakout the fiber channel ports, the admin status is disabled by default. Enable the status if required.

To switch the admin status:

Select the port and click **Toggle Admin State**.

The **Toggle Admin State** window is displayed.

Configuring Maximum Transmission Unit

You can configure the Maximum Transmission Unit (MTU) for full-switch and fabric mode IOMs.

To configure MTU:

1. Click **Devices > I/O Modules > View Details > Hardware > Port Information**.

2. Select the Ethernet port and click **MTU**.
The **Configure MTU** window is displayed.

3. Select the **MTU Size**.

The approximate value for MTU is 1500 bytes. The default value is 1532 bytes, and the maximum is 9000 bytes. If the port has both FCoE and Ethernet, the value is 2500 bytes.

Configuring auto negotiation

You can switch auto negotiation (AutoNeg). For DAC cabling, the AutoNeg is enabled by default. For AOC (fiber), AutoNeg is disabled by default.

To switch AutoNeg:

Select the port and click **Toggle AutoNeg**.

The **Toggle AutoNeg** window is displayed.

If Ethernet links are not displayed automatically, switch the auto negotiation setting.

MX Scalable Fabric architecture

The scalable fabric architecture ties multiple MX7000 chassis into a single network domain to behave like a single logical chassis from a networking perspective. The MX scalable fabric architecture provides multi-chassis Ethernet with:

- Multiple 25Gb Ethernet connections to each server sled
- No east-west oversubscription
- Low “any-any” latency
- Scale up to 10 MX7000 chassis
- Flexible uplink speeds
- Support for non-PowerEdge MX devices such as rack servers

For more information, see the *PowerEdge MX I/O Guide* available at www.dell.com.

Architectural Overview

A scalable fabric consists of two main components – a pair of MX9116n Fabric Switching Engines (FSE) and additional pairs of MX7116n Fabric Expander Modules (FEM) used to connect remote chassis to the FSEs. This is a hardware enabled architecture and it applies irrespective of whether the switch is running in Full Switch or Fabric modes. A total of ten MX7000 chassis are supported in a scalable fabric.

Fabric Switching Engine

The FSE contains the switching ASIC and network operating system. Traffic that is received from a FEM is mapped to the correct switch interface automatically. Each NIC port has a dedicated 25GbE lane from the NIC through the FEM and into the FSE so there is no port to port oversubscription.

Fabric Expander Module

An FEM takes Ethernet frames from a compute node and sends them to the FSE and from the FSE to the compute node. There is no switching ASIC or operating system running on the FEM, which allows for a very low latency. This also means that there is no firmware that needs to be updated. The FEM is invisible to the FSE and does not need to be managed in any way.

When using dual-port NICs, only the first port on the FEM must be connected to the FSE. The second port is not used.

When connecting a FEM to an FSE, the rules to remember are:

- FEM in Slot A1 connects to FSE in Slot A1
- FEM in Slot A2 connects to FSE in Slot A2
- FEM in Slot B1 connects to FSE in Slot B1
- FEM in Slot B2 connects to FSE in Slot B2

Topics:

- [Recommended physical topology](#)
- [Restrictions and guidelines](#)
- [Recommended connection order](#)

Recommended physical topology

The recommended minimum design for a scalable fabric is two chassis with fabric A populated with redundant IOMs. Ideally, the two chassis are located in separate racks on separate power circuits to provide the highest redundancy.

Additional chassis only have FEMs and they appear as the image below.




Table 8. Fabric topology

Chassis	Slot	Module
Chassis 1	A1	MX9116n FSE
	A2	MX7116n FEM
Chassis 2	A1	MX7116n FEM
	A2	MX9116n FSE
Chassis 3-10	A1	MX7116n FEM
	A2	MX7116n FEM

You can also use Fabric B to create a second scalable fabric:



Restrictions and guidelines

The following restrictions and guidelines are applicable when building a scalable fabric:

- Mixing switch types in the same fabric is not supported. For example: MX9116n in slot A1 and MX5108n in slot A2
- Mixing switch types across fabrics is supported. For example: MX9116n in slots A1& A2 and MX5108n in slots B1 & B2
- All FSE and FEM IOMs in a scalable fabric must be in the same OME–Modular MCM group. FEMs in a chassis in MCM group 1 cannot be connected to FSEs in a chassis in MCM group 2.

The following restrictions are applicable when implementing a scalable fabric in both fabric slot A and fabric slot B:

- IOM placement for each scalable fabric must be the same within the same chassis. For example, if the FSE for the first scalable fabric is in Slot A1, then the second FSE must be in slot B1 in the same chassis, and so on.

- For chassis that only contain FEMs, all four FEMs must connect to the same chassis with the FSEs. The fabric B FEMs cannot be connected to FSEs in a different chassis as the fabric A FSEs.

Recommended connection order

Any QSSFP28-DD port on the MX9116n can be used for any purpose. The table below describes the recommended port order for connecting chassis with Fabric Expander Modules (FEMs) to the FSE. The table contains references IOMs in fabric A, but the same guidelines apply to IOMs in fabric B.

Table 9. Recommended port order for connecting FEM to FSE

Chassis	FSE Port (Phys Port)
1 & 2	FSE Port 1 (17/18)
3	FSE Port 7 (29/30)
4	FSE Port 2 (19/20)
5	FSE Port 8 (31/32)
6	FSE Port 3 (21/22)
7	FSE Port 9 (33/34)
8	FSE Port 4 (23/24)
9	FSE Port 10 (35/36)
10*	FSE Port 6 (25/26)

*—By default, the port group 10 is not configured to support a FEM. If you want to connect a FEM to this port, use the OME - Modular interface to set the port mode to Fabric Expander.



NOTE: The port groups, 6, 11, and 12 (physical ports 27/28, 37/38, 39/40), can be used for additional uplinks, ISLs, rack servers, and so on.

SmartFabric Services

SmartFabric Services is a capability of Dell EMC Networking OS10 Enterprise Edition running on Ethernet switches designed for the PowerEdge MX platform.

A SmartFabric is a logical entity containing a collection of physical resources such as servers and switches and logical resources—networks, templates, and uplinks. In the SmartFabric Services mode, the switches operate as a simple Layer 2 input output aggregation device, which enables complete interoperability with network equipment vendors.

A SmartFabric provides:

- Data center Modernization
 - I/O Aggregation
 - Plug-and-play fabric deployment
 - A single interface to manage all switches in the fabric like a single logical switch
- Lifecycle Management
 - Fabric-wide firmware upgrade scheduling
 - Automated or user enforced roll back to last well-known state
- Fabric Automation
 - Ensured compliance with selected physical topology
 - Policy based Quality of Service (QoS) based on VLAN and Priority assignments
 - Automatic detection of fabric misconfigurations and link level failure conditions
 - Automated healing of the fabric on failure condition removal
- Failure Remediation
 - Dynamically adjusts bandwidth across all inter-switch links in the event of a link failure

Unlike Full Switch mode, most fabric configuration settings are performed using the OME-Modular.

For information about automated QoS, see [SmartFabric VLAN management and automated QoS](#)

Changing operating modes

In both Full Switch and Fabric modes, all configuration changes you make using the OME-Modular interface are retained when you switch modes. It is recommended that you use the GUI for all switch configurations in Fabric mode and the OS10 CLI for configuring switches in Full Switch mode.

To switch an MX9116n Fabric Switching Engine or MX5108n Ethernet Switch between Full Switch and Fabric modes, use the OME-Modular GUI and create a fabric with that switch. When that switch is added to the fabric, it automatically changes to Fabric mode. When you change from Full Switch to Fabric mode, all Full Switch CLI configuration changes are deleted except for a subset of settings supported in Fabric mode.

To change a switch from Fabric to Full Switch mode, the fabric must be deleted. At that time, all Fabric GUI configuration settings are deleted. However, the configurations supported by the subset of Fabric CLI commands (hostname, SNMP settings, and so on) and the changes you make to port interfaces, MTU, speed, and auto-negotiation mode, are not deleted. The changes to port interfaces exclude the administrator state—shutdown/no shutdown.

Topics:

- [Guidelines for operating in SmartFabric mode](#)
- [SmartFabric network topologies](#)
- [Switch to switch cabling](#)
- [Upstream network switch requirements](#)
- [NIC teaming restrictions](#)
- [CLI commands available in Fabric mode](#)
- [Viewing fabric details](#)
- [Adding fabric](#)
- [Deleting fabric](#)

- [VLANs for SmartFabrics and FCoE](#)

Guidelines for operating in SmartFabric mode

The guidelines and restrictions while operating in SmartFabric mode are as follows:

- When operating with multiple chassis, ensure that the switches in A1/A2 or B1/B2 in one chassis are interconnected only with other A1/A2 or B1/B2 switches respectively. Connecting switches that are placed in slots A1/A2 in one chassis with switches in slots B1/B2 in another chassis is not supported.
- Uplinks must be symmetrical. If one switch in a SmartFabric has two uplinks, the other switch must have two uplinks of the same speed.
- Enable LACP on the uplink ports for switches being uplinked.
- You cannot have a pair of switches in SmartFabric mode uplink to another pair of switches in SmartFabric mode. You can only uplink a SmartFabric to a pair of switches in Full Switch mode.


SmartFabric network topologies

The SmartFabric Services support three network topologies with specific IOM placement requirements.

- 2 x MX9116n Fabric Switching Engines in different chassis
- 2 x MX5108n Ethernet Switches in the same chassis
- 2 x MX9116n Fabric Switching Engines in the same chassis


2 x MX9116n Fabric Switching Engines in separate chassis

This placement is recommended while creating a SmartFabric on top of a scalable fabric architecture. This configuration supports placement in Chassis1/A1 and Chassis 2/A2 or Chassis1/B1 and Chassis 2/B2. A SmartFabric cannot include a switch in Fab A and a switch in Fab B. If one of the chassis fails, placing the FSE modules in separate chassis provides redundancy. Both the chassis must be in the same MCM group.



2 x MX5108n Ethernet Switches in the same chassis


The MX5108n Ethernet Switch is supported only in single chassis configurations. The switches must be placed in slots A1/A2 or slots B1/B2. A SmartFabric cannot include a switch in Fab A and a switch in Fab B.



In SmartFabric mode, ports 9 and 10 are automatically configured in a VLT at 40GbE speed. For port 10, use a cable or optic that supports 40GbE and not 100GbE.

2 x MX9116n Fabric Switching Engines in the same chassis

Use this placement in environments with a single chassis. The switches must be placed in either slots A1/A2 or slots B1/B2. A SmartFabric cannot include a switch in Fab A and a switch in Fab B.




The fabric design, "2 x Mx9116n Fabric Switching Engines in the same chassis" is supported, but not recommended. Use of this design displays an error message on the **Fabric Topology** and **View Topology** pages of OME - Modular.

Switch to switch cabling

When operating in SmartFabric mode, each switch pair runs a Virtual Link Trunk (VLT) link between them. For the MX9116n, the port groups 11 and 12 are used.



For the MX5108n, ports 9 and 10 are used. Port 10 operates at 40GbE instead of 100GbE because all VLT links must run at the same speed. Ensure that you use a cable or optic fibre that supports 40GbE.



NOTE: You cannot select the ports, and the connection topology is enforced by SmartFabric Services.

NOTE: VLT is supported only on Ethernet and not on FCoE. Physically separate uplinks for LAN and FCoE traffic are required for MX5108n and MX9116n switches.

Upstream network switch requirements

It is recommended, but not required, that PowerEdge MX switches are connected to a pair of redundant upstream switches. When you are connecting a pair of switches in Fabric mode to an upstream switch pair, the upstream switch pair **must** be configured in the following way:

- Both upstream switches must be connected to each other using technologies such as VLT or VPC.
- The upstream switch ports must be in a port channel using LACP.
 - NOTE:** The LACP option is supported on Ethernet uplinks only.
- Ensure that a compatible Spanning Tree Protocol is configured. For more information, see the section, **Spanning Tree Protocol**.

Spanning Tree Protocol

OS10 defaults to RPVST+ as the Spanning Tree protocol. To change STP modes, use the spanning-tree mode command. Use the spanning-tree mode command to change STP modes. For steps, see the *OS10 Enterprise Edition User Guide*.

NOTE: If the upstream network is running RSTP, change from RPVST+ to RSTP before physically connecting the switches to the upstream network. Failure to do so may cause a network outage.

NIC teaming restrictions

NIC teaming is suggested for redundancy unless a particular implementation recommends against it. There are two main kinds of NIC teaming:

- Switch Dependent—Also referred to as 802.3ad or Dynamic Link Aggregation. The switch dependent teaming method uses the LACP protocol to understand the teaming topology. This teaming method provides Active-Active teaming and requires the switch to support LACP teaming.
- Switch Independent—This method uses the operating system and NIC device drivers on the server to team the NICs. Each NIC vendor may provide slightly different implementations with different pros and cons.

NIC Partitioning (NPAR) can impact how NIC teaming operates. Based on restrictions implemented by NIC vendors related to NIC partitioning, certain configurations preclude certain types of teaming.

The following restrictions are applicable to both Full Switch and SmartFabric modes:

- If NPAR is NOT in use, both Switch Dependent (LACP) and Switch Independent teaming methods are supported.
- If NPAR IS in use, only Switch Independent teaming methods are supported. Switch Dependent teaming is NOT supported.

The following restrictions are applicable to Switch Dependent (LACP) teaming:

- The IDRAC shared LOM feature can only be used if “Failover” option on IDRAC is enabled
- If the host OS is Windows, the LACP timer MUST be set to “slow” (also referred to as “normal”)

For the list of supported operating systems, see *Dell EMC PowerEdge MX7000 Enclosure Installation and Service Manual*.

NOTE: In the fabric mode, if an LACP team is created with four port and you want to delete two ports from the LACP team, you must delete the entire LACP team and create a new LACP team with two ports.

For detailed NIC teaming instructions, refer to the network adapter or operating system documentation.

CLI commands available in Fabric mode

When operating in Fabric mode, most of the switch configuration is managed through the OME–Modular GUI. Some OS10 functionality, such as Layer 3 routing, are disabled. Owing to the disabling, a switch operating in Fabric mode supports all OS10 show commands, but only the below subset of CLI configuration commands:

- `clock`—Configure clock parameters.
- `end`—Exit to the EXEC mode.
- `exit`—Exit from current mode
- `help`—Display available commands.
- `hostname`—Set the system hostname.
- `interface`—Configure or select an interface.
- `ip name-server`—Configure the IP address of up to three name servers
- `logging`—Configure system logging.
- `management route`—Configure the IPv4/IPv6 management route
- `no`—Delete or disable commands in Configuration mode.
- `ntp`—Configure Network time protocol.
- `snmp-server`—Configure SNMP server.
- `username`—Create or modify user credentials.
- `spanning-tree`
 - `disable`— Disable spanning tree globally.
 - `mac-flush-timer`— Set the time used to flush MAC address entries.
 - `mode`— Enable a spanning-tree mode, such as RSTP or MST.
 - `mst`— Configure multiple spanning-tree (MST) modes.
 - `rstp`— Configure rapid spanning-tree protocol (RSTP) mode.
 - `vlan`— Configure spanning-tree on a VLAN range.
 - `username`— Create or modify user credentials.
 - `SupportAssist`— Configure SupportAssist related settings.
 - `Security`— Configure features corresponding to network security.
 - `Fibre Channel`— Configure features corresponding to fibre channel interfaces.

Viewing fabric details

To view details an existing fabric:

- From the **Devices** drop down, select **Fabric**.
- From the fabrics table, select the fabric and click **View Details**.

The **Fabric Details** page is displayed.

Adding fabric

To add a fabric:

1. Click **Devices > Fabric** .
The **Fabric** page is displayed.
2. Click **Add Fabric**.
The **Create Fabric** window is displayed.
3. Enter **Name** and **Description**, and then click **Next**.
4. Select the **Design Type** from the drop-down.
The available options are:
 - 2xMX5108n Ethernet Switches in same chassis

- 2xMX9116n Fabric Switching Engines in same chassis
- 2xMX9116n Fabric Switching Engines in different chassis

Based on the design type selected, the options to select the chassis and the switches—A and B, are displayed.

5. Select the chassis and switches.
The cabling image is displayed.
6. Click **Next** to view the summary of the fabric.

You can print to print a hard copy of the fabric details or save the details as a PDF on your system.

After the fabric is created, the switch is placed in the SmartFabric mode and the IOM reboots.

 **NOTE: After a fabric is created, the health status of the fabric is critical until uplinks are created.**

 **NOTE: The fabric health alerts are displayed on all chassis in the MCM group.**

Adding uplinks


To add uplinks:

1. From the **Devices** drop-down, select **Fabric**.
The **Fabric** page is displayed.
2. From the fabrics table, select the fabric and click **View Details**.
The **Fabric Details** page is displayed.
3. From the **Uplinks** section, click **Add Uplink**.
The **Add Uplink** window is displayed.
4. Enter **Name**, **Description**, select the **Uplink Type**, and then click **Next**.

The available options are:

- **Ethernet**—You can pick one or more Ethernet ports across switches to form a LAG. The network can be of any type. For example: Ethernet.
- **FCoE**—You can pick one port from an IOM and associate a single network of FCoE type. This is for FCoE connectivity that connects to another switch that connects to the FC network. For single fabric, you can have two FCoE uplink, one from each IOM. Both IOMs must have different network that is, different FCoE VLANs.

In FCoE mode, untagged VLAN on the server port and FCoE uplink must be the same. This condition ensures that the untagged FIP VLAN discovery (L2 frame) packets are switched to the untagged VLAN. The FCoE uplink is used to identify FIP Snooping Bridge (FSB) mode at the switch. For the FCoE sessions to come up, configure the same untagged VLAN on FCoE uplinks and server ports.

 **NOTE: On the uplink FCoE switch, use the default fc-map (0efc00) only.**

- **FC Gateway**—You can pick one or more ports from the same IOM and associate a single network of FCoE type. This type of uplink is for connectivity to a SAN switch. For single Fabric, you can have two FC gateway uplinks one from each IOM. Both IOMs must have different network that is, different FCoE VLANs. For a given fabric, you can have at least one uplink of type FC (either of FCoE, FCDirectAttach, FC Gateway).

In Fabric mode, you can assign any untagged VLAN to Ethernet server ports that belong to a FCoE VLAN that has one or more FC Gateway uplinks. The FC Gateway uplink is used to identify NPG (N Port Proxy Gateway) mode at the switch.

- **FC Direct Attach**—You can pick one or more ports from same IOM and associate a single network of FCoE type. This type of uplink is for direct FC storage connectivity. For single fabric, user can have two FC DirectAttach uplink, one from each IOM. Both IOMs must have different networks that is, different FCoE VLANs.

In Fabric mode, you can assign any untagged VLAN to Ethernet server ports that belong to a FCoE VLAN that has one or more FC Direct attach uplinks. The FC Direct attach uplink is used to identify F-Port mode at the switch.

5. Choose the necessary **Switch Ports** and select any **Tagged Networks** .

If you are required to configure new network other than the existing ones, click **Add Network** and enter the network details. For more details see, [Adding Network](#).

Adding network

You can use the **Fabric** and **Configuration > Network** pages to add networks. For more information, see [Defining networks](#).

To add a new network from the **Fabric** page:

1. From the **Devices** drop down, select **Fabric**.
The **Fabric** page is displayed.
2. From the fabrics table, select the fabric and click **View Details**.
The **Fabric Details** page is displayed.
3. From the **Uplinks** section, click **Add Uplink**.
The **Add Uplink** window is displayed.
4. Click **Add Network**.
The **Define Network** window is displayed.
5. Enter **Name**, **Description**, **VLAN ID** and select the **Network Type**.
For the network types, see the *Online Help*.

Editing uplink

To edit an existing uplink:

1. From the **Devices** drop-down, select **Fabric**.
The **Fabric** page is displayed.
2. From the fabrics table, select the fabric and click **View Details**.
The **Fabric Details** page is displayed.
3. From the **Uplinks** table, select the uplink and click **Edit**.
The **Edit uplink** page is displayed.
4. Edit the **Name**, **Description**, and **Uplink Type** fields as necessary, and then click **Next**.
5. Select the necessary **Switch Ports** and select any **Tagged Networks** or **Untagged Networks**.
To configure new network other than the existing ones, click **Add Network** and enter the network details. For more details see, [Adding Network](#).

 **NOTE:** You cannot edit the ports or networks when uplinks are in FCoE, FC Gateway, or FC Direct Attach modes.

Viewing topology details

The fabric topology image displays only the operational status of the ports. If the operational status is "up", a check mark is displayed. To view the graphical representation of the validation errors in an MCM scenario, go to the **Group Topology** page on the OME-Modular web interface.

To view topology details:

- From the **Devices** drop-down, select **Fabric**.
- From the fabrics table, select the fabric and click **View Details**.
- From the **Fabric Details** page, click **Topology**.

The topology of the fabric is displayed.

Editing fabric details

To edit the fabric details:

1. From the **Devices** drop down, select **Fabric**.
The **Fabric** page is displayed.
2. From the fabrics table, select the fabric and click **Edit**.
The **Edit Fabric** page is displayed.
3. Make the necessary changes to the **Name** and **Description** fields.

Deleting uplinks

To delete an uplink:

1. From the **Devices** drop down, select **Fabric**.
The **Fabric** page is displayed.
2. In the fabrics table, select any fabric and click **View Details**.
3. In the uplinks table, select the uplink to be deleted.

4. Click **Delete**. Click **Yes** to confirm the deletion.

Deleting fabric

To delete an existing fabric:

1. From the **Devices** drop-down, select **Fabric**.
The **Fabric** page is displayed.
2. From the fabrics table, select the fabric that you want to delete.
3. Click **Delete**.
A message is displayed prompting you to confirm the deletion.
4. Click **Yes** to proceed.
After the fabric is deleted, the IOM reboots.

VLANs for SmartFabrics and FCoE

Create VLANs before creating the SmartFabric. The first VLAN that is created must be the default or native VLAN, typically VLAN 1. The default VLAN must be created for any untagged traffic to cross the fabric.

If you are implementing Fibre Channel configurations, you can also configure VLANs for FCoE. The storage arrays have two separate controllers that create two paths—SAN path A and SAN path B. These paths are connected to MX9116n FSE. For storage traffic to be redundant, two separate VLANs are created for that traffic.

The following table lists examples of VLAN attributes for FCoE traffic:

Table 10. VLAN attributes for FCoE

Name	Description	Network Type	VLAN ID	SAN
FC A1	FCOE A1	Storage - FCoE	30	A
FC A2	FCOE A2	Storage - FCoE	40	B

Defining VLANs

To define VLANs:

1. From the menu, click **Configuration > Networks**.
2. In the **Network** pane, click **Define**.
The **Define Network** window is displayed.
3. Enter a **Name** and **Description** for the VLAN.
The description is optional.
4. Enter a **VLAN ID** and select the **Network Type**.
For FCoE, the **Network Type** must be **Storage FCoE**.
5. Click **Finish**.

Editing VLANs

You can add or remove VLANs on the deployed servers in a SmartFabric.

To add or remove VLANs:

1. From the menu, click **Devices > Fabric**.
2. Select the fabric for which you want to add or remove the VLAN.
3. In the left pane, select **Servers** and select the required servers.
4. Click **Edit Networks**.
5. Select one of the following options:
 - **NIC teaming from LACP**
 - **No Teaming**
 - **Other**

6. Define the tagged and untagged VLANs to modify the VLAN selections as required.
7. Select VLANs on Tagged and Untagged Network for each Mezzanine card port.
8. Click **Save**.

VLAN scaling guidelines

The number of recommended VLANs differs between the modes as SmartFabric mode provides network automation capabilities that Full Switch mode does not.

The following table lists the maximum number of VLANs recommended per fabric, Uplink, and server port:

Table 11. Maximum number of VLANs recommended in SmartFabric mode

OS10 Version	Parameter	Value
10.5.0	Maximum VLANs per fabric	256
	Maximum VLANs per uplink	256
	Maximum VLANs per server port	64
10.4.0.R3S	Maximum VLANs per fabric	128
10.4.0.R4S	Maximum VLANs per uplink	128
	Maximum VLANs per server port	32

Managing networks

You can configure logical networks that represent your environment, for the tagged and untagged VLANs. These logical networks are used to provision the appropriate VLANs on the associated switch port for the physical server NIC port.

NOTE: VLANs are only assigned to servers connected to switches in Fabric mode. For servers connected to switches in Full Switch mode, the VLAN information is ignored.

In tagged networks, a port handles multiple VLANs. VLAN tagged networks help identify which packet belongs to the VLAN on the other side. A packet is tagged with a VLAN tag in the Ethernet frame. A VLAN ID is put in the header to identify the network to which it belongs.

In untagged networks, one port handles only one VLAN.

To view the list of networks, click **Configuration > Networks**. The **Networks** page with the list of networks is displayed. You can view the name, description, and VLAN ID of the networks.

A summary of the selected network is displayed on the right side.

You can perform the following tasks on the **Networks** page:

- Define networks
- Edit networks
- Delete networks
- Export networks

Topics:

- [SmartFabric VLAN management and automated QoS](#)
- [Defining networks](#)
- [Editing networks](#)
- [Exporting network configurations](#)
- [Deleting network configurations](#)

SmartFabric VLAN management and automated QoS

Besides assigning VLANs to server profiles, SmartFabric Services automate QoS settings based on user input. When a VLAN is created and you select the related traffic type (such as iSCSI and vMotion), the SFS engine assigns the correct QoS setting to that VLAN. You can also select a “metal” such as gold and bronze to assign your own priority values to the traffic.

Table 12. Network traffic types - QoS settings

Network Traffic Type	Description	QoS Setting
General Purpose (Bronze)	Used for low-priority data traffic	2
General Purpose (Silver)	Used for standard/default priority data traffic	3
General Purpose (Gold)	Used for high-priority data traffic	4
General Purpose (Platinum)	Used for extremely high-priority data traffic	5
Cluster Interconnect	Used for cluster heartbeat VLANs	5
Hypervisor Management	Used for hypervisor management connections such as the ESXi management VLAN	5
Storage - iSCSI	Used for iSCSI VLANs	5

Table 12. Network traffic types - QoS settings(continued)

Network Traffic Type	Description	QoS Setting
Storage - FCoE	Used for FCoE VLANs	5
Storage - Data Replication	Used for VLANs supporting storage data replication such as for VMware VSAN	5
VM Migration	Used for VLANs supporting vMotion and similar technologies	5
VMWare FT Logging	Used for VLANs supporting VMware Fault Tolerance	5

Defining networks

To configure a logical network:

1. Click **Configuration > Networks**.
The **Networks** page is displayed.
2. Click **Define**.
The **Define Network** window is displayed.
3. Enter the name, description, VLAN ID.
The format for a single VLAN ID is—123 while for an ID range, the format is—123-234.
4. Select the **Network Type**.
For more details, see [SmartFabric VLAN management and automated QoS](#)The available options are:
 - **General Purpose (Bronze)**
 - **General Purpose (Silver)**
 - **General Purpose (Gold)**
 - **General Purpose (Platinum)**
 - **Cluster Interconnect**
 - **Hypervisor Management**
 - **Storage - iSCSI**
 - **Storage - FCoE**
 - **Storage - Data Replication**
 - **VM Migration**
 - **VMWare FT Logging**

For more details, see [SmartFabric VLAN management and automated QoS](#).

Editing networks

To edit a network:

1. On the **Networks** page, select the network that you want to edit, and click **Edit**.
The **Edit Network** window is displayed.
2. Make the required changes.
While editing the network, ensure that only one VLAN is configured in both the ports.

 **NOTE: In fabric mode, do not delete VLAN from OME–Modular, if the VLAN is associated with any uplink.**

Exporting network configurations

To export the network configuration:

On the **Networks** page, select the desired network and click **Export**.
The network details are exported in a .csv format to a local drive on your system.

Deleting network configurations

To delete a network:

On the **Networks** page, select the network and click **Delete**.

If the network is associated with a fabric uplink, a warning message is displayed that deleting the network results in loss of connectivity.

Managing Fibre Channel IOMs

The MXG610s Fibre Channel (FC) switch is designed for mission critical applications accessing data on external storage. It is optimized for flash storage and virtualized server environments. The FC switch enables organizations to dynamically scale connectivity and bandwidth Ports-on-Demand (PoD). It enhances operations with consolidated management and simple server and storage connectivity.

OME–Modular makes the management of the MXG610s simple. The SSO feature in OME–Modular enhances security and convenience.

To view GUI of the MXG610s FC switch:

1. On the **Devices > I/O Modules >** page, click **IOM UI launch**.

The MXG610s FC web tools interface is displayed.

Managing firmware

The firmware feature in OME–Modular helps you to update the firmware of all the components in the chassis. The components include compute sleds, ethernet IOMs, storage IOMs, and SAS IOMs. The firmware updates can be sources from the Dell web site or a custom repository setup using Repository Manager.

You must have the chassis administrator role and the device update privilege for the chassis to update the firmware on the chassis. To update the firmware on the components, you must have the device-specific manager role and device update privilege to perform the updates.

The MX chassis bundle refers to the following update packages:

- Chassis manager DUP—This DUP comprises of the OME–Modular firmware.
- Storage sled DUP—This DUP contains updates for the Dell storage sleds in the chassis.
- Storage IOM DUP—This DUP contains updates for the chassis storage IOMs.

The DUPs for network IOMs and switches are licensed software and are available as individual DUPs. For external storage, the DUPs are bundled in the catalog. If the hard drives or storage enclosures are assigned to a compute sled, you can update them using iDRAC. However, you cannot update the assigned or unassigned hard drives through a chassis context. You can map the drives to a server to update them.

The compute sled bundle refers to the packages for the server components—BIOS, NIC, RAID, hard drives, and iDRAC.

The firmware update process involves specifying the catalog, retrieving the firmware inventory, checking compliance, and updating the firmware.

The available baselines are displayed on the **Configuration > Firmware** page. You can view a summary of the baseline compliance and a pie chart on the top of the page. You can also view the summary of the desired baseline on the right side of the **Firmware** page.

The baseline information that is displayed on the **Firmware** page is—compliance, name of the baseline, job status, catalog type, timestamp when the baseline was last used.

You can perform the following tasks on the **Firmware** page:

- Create baseline
- Edit baseline
- View report
- Delete baseline
- Manage catalogs
- Check compliance

Topics:


- [Creating baselines](#)
- [Checking compliance](#)
- [Editing baselines](#)
- [Managing catalogs](#)
- [Updating firmware](#)
- [Rolling back firmware](#)
- [Deleting firmware](#)

Creating baselines

To create a firmware baseline:

1. Click **Configuration > Firmware > Create Baseline** .
The **Create Firmware Baseline** window is displayed.
2. Select the catalog type, enter a name and description for the baseline.
3. Click **Add**.
The **Add Firmware Catalog** window is displayed.
4. Select the catalog source.

5. In the **Create Firmware Baseline** window, select the devices and groups for which you want to create the baseline. After the baseline is created, a message is displayed and a compliance check is performed on the baseline. The status of the job is displayed on the **Firmware** page.

 **NOTE:** If the baseline is created from the catalog, the information of the associated baseline is displayed.

Checking compliance

To check the compliance of a firmware baseline:

1. On the **Firmware** page, select the baseline and click **Check Compliance**.
A summary of the compliance check is displayed on the right side of the **Firmware** page.
2. Click **View Report**.
The **Compliance Report** page is displayed.

You can view details including the name of the catalog and baseline, status of the compliance, type of the baseline, name of the device, model, service tag of the device, current update version, and baseline version.

You can perform the following tasks on the **Compliance Report** page:

- Update firmware
- Export the report in .csv format to a local drive on your system.
- Sort the device information using **Advanced Filters**

When you update the firmware for SAS IOMs that are available as an individual component and a chassis component, using the compliance report method, the management module update fails. Select the SAS IOM from the chassis component or the SAS IOM listed individually in the compliance report.

Editing baselines

To edit a baseline:

1. On the **Firmware** page, select the baseline that you want to modify and click **Edit**.
The **Edit Firmware Baseline** window is displayed.
2. Make the required changes.

Managing catalogs

The catalog management feature in OME–Modular helps you to configure the catalog location and create firmware baselines. A catalog contains metadata of the bundles and individual DUPs or packages. The bundles represent package sets that are tested and certified together.


The catalogs can be sourced from the following locations:

- Dell website—You can specify the proxy parameters to enable the application to access the Internet from your network. The proxy parameters include network address and optional credentials—user name and password. The proxy settings are configured during initial setup or on the **Application Settings > Network** page.

Multiple catalogs could be posted on the Dell website.

- Network share or website location in your network—The network share consists of NFS, CIFS, HTTP, or HTTPS.

You can use the Repository Manager to create the catalog and store it on the network share. If you have the chassis administrator privilege, you can view the list of catalogs and perform basic management tasks such as editing and deleting the catalogs. You cannot delete a catalog that is associated with a baseline. If a catalog is inaccessible, an operational status icon is displayed for the catalog.

 **NOTE:** When you create a catalog on a particular date and download it to the required location on your network or local drive, the download is successful. However, if you modify the catalog on the same day at different times and attempt downloading it, the modified catalog is not downloaded. If the repository type is NFS and the catalog file is not available on the specified NFS server, the system uses the catalog file that was last fetched.

To view the list of catalogs:

- On the **Firmware** page, click **Catalog Management**.
The **Catalog Management** page is displayed.

You can select a catalog to view the summary on the right-side. The summary consists of the number of bundles in the catalog, date and time when the catalog was released, and name of the baselines associated with the catalog.

You can perform the following tasks on the **Catalog Management** page:

- Add catalogs
- Edit catalogs
- Delete catalogs

Viewing catalogs

You can view the following catalog information on the **Catalog Management** page.

- Name and download status of the catalog
- Type of the repository from where the catalog is downloaded
- Location of the repository
- Name of the catalog .xml file
- Release timestamp of the catalog

1. On the menu bar, click **Configuration > Firmware > Catalog Management**.

The **Catalog Management** page is displayed.

2. Select a catalog to view the summary on the right side.

The summary comprises of the number of bundles in the catalog, release timestamp of the catalog, and the name of the associated bundles in the catalog.

Adding catalogs

To add catalogs:

1. On the **Catalog Management** page, click **Add**.

The **Add Firmware Catalog** window is displayed.

2. Enter a name for the catalog and select the catalog source.

The available options are:

- **Newest validated stacks of chassis firmware on Dell.com**—The versions of firmware in this catalog have been tested together as part of the latest OME - Modular firmware release.
- **Latest component firmware versions on Dell.com**—This catalog may include versions of firmware for components that have been individually released since the last validated stack of chassis firmware.
- **Network Path**—The folder where a catalog and optionally associated updates have been placed by unpacking the validated stack at **ftp.dell.com** or by using Dell EMC Repository Manager.

3. Select the **Share Type**.

The available options are:

- NFS
- CIFS
- HTTP
- HTTPS

NOTE: The **Share Type** option is available only if you select **Network Path**.

NOTE: The **HTTPS share feature with proxy does not work when authentication is enabled for both the proxy and HTTPS share**.

4. Select the mode of updating the catalog.

The available options are:

- Manually
- Automatically

The default mode is manual.

5. Select the **Update Frequency**.

- Daily

- Weekly

The time can be in HH:MM format.

Editing catalogs

You can only modify the catalog name, network share address, and catalog filepath.

To edit catalogs:

1. On the **Catalog Management** page, select the catalog that you want to edit and click **Edit**. The **Edit Firmware Catalog** window is displayed.
2. Make the required changes.

Deleting catalogs

You can only delete catalogs that are not associated with a baseline. If you attempt deleting a catalog that is associated with a baseline, an error message is displayed.

To delete a catalog:

On the **Catalog Management** page, select the catalog that you want to delete and click **Delete**.

Updating firmware

Before updating the firmware on a chassis, compute, or storage sleds, ensure that all IOMs and network fabrics are healthy. You can upgrade only six IOMs simultaneously.

NOTE: The **Update Firmware** button may be disabled temporarily during inventory refresh when a **Refresh Inventory job** or **Default Inventory job** is run.

To update firmware:

1. On the **Compliance Report** page, select the device or component for which you want to update the firmware. The **Update Firmware** window is displayed.
2. Select the **Update Now** option to update the firmware immediately or **Schedule Later** to update the firmware on the chosen date and time.

NOTE: If the system displays the local clock on the **Time Configuration** page even after you configured the NTP servers, reconfigure the NTP servers.

NOTE: During firmware update, when the active MM reboots and the standby MM is active, some messages on the **Execution Details** page for the firmware update are not displayed. The messages are not displayed owing to synchronization issues.

NOTE: During the OME–Modular firmware update, multiple users can upload the OME–Modular DUP using any interface. However, a warning message may be displayed after the firmware update job is initiated.

NOTE: For non-default VLAN, the management IPv6 IP of MX9116n or MX5108n IOMs is unreachable if, the DHCP V6 configuration in ToR switch does not have the IPV6 default gateway.

Rolling back firmware

If you are not convinced with the firmware update of a device or component, you can roll back the update to the version before the update. The rollback option is enabled only if OME–Modular can access the firmware package of the previous version. The following methods can be used to enable the access:

- A Device that has the rollback version (or N-1 version) that matches the previous version. Not all devices support a rollback or N-1 version. The rollback version is displayed as a rollback candidate even if it does not match the version before the update.
- An imported catalog that has a reference to the previous catalog version.
- You can browse for a firmware package that has the previous firmware version.

For Network IOMs, the availability of rollback information depends on the status of the Network IOM (Full Switch or Fabric) and the firmware update method. If the firmware is updated on nodes in the fabric, the rollback information is available on the node on which the

firmware update is initiated. If the firmware on the member chassis Network IOMs is updated through the Lead chassis, the rollback information is available on only on the Lead chassis.

To roll back a firmware update:

1. On the **Firmware** page, click **Rollback Firmware**.
The **Rollback Firmware** window is displayed.
2. Select the component for which you want to roll back the firmware and click **Rollback**.

Deleting firmware

You can delete firmware baselines, if you have the administrator privilege.

To delete a firmware baseline:

On the **Firmware** page, select the baseline that you want to delete, and click **Delete**.
A message is displayed prompting you to confirm the delete operation.

Monitoring alerts and logs

You can view and manage the alerts that are generated in the management system environment. You can filter alerts and perform the appropriate actions.

Every chassis in the MCM group receives Fabric alerts, irrespective of whether the MX5108N or MX9116N IOMs present in the chassis to accommodate new MX5108N or MX9116N IOMs in the chassis.

To view the alerts page, from the menu bar, click **Alerts**. The **Alerts** page with the following tabs is displayed:

- **Alert Log**
- **Alert Policies**
- **Alert Definition**

Topics:

- [Alert log](#)
- [Alert policies](#)
- [Alert definitions](#)

Alert log

The **Alerts Log** page displays the list alert logs for events occurring in the chassis. On the menu bar, click **Alerts > Alert Log**. The **Alerts Log** page is displayed. You can view the alerts details—severity of the alert, timestamp, source, category, subcategory, message ID, and description of the alert.

The **Alerts Log** page displays 30,000 records. Select an alert to view the summary of the alert on the right side of the **Alerts Log** page. You can also perform the following tasks on the **Alerts Log** page:

- Acknowledge alerts
- Unacknowledge alerts
- Ignore alerts
- Export alerts
- Delete alerts

The latest unacknowledged alerts are displayed on the OME–Modular home page.

Filtering alert logs

To filter alert logs:

1. On OME–Modular web interface, navigate to **Alerts > Alert Log**.
2. Click **Advanced Filters**.
3. Select or update the following based on your requirement:
 - **Severity**—To view all alerts with specific severity level.
 - **Acknowledge**—To view all alerts that were acknowledged.
 - **Start Date** and **End Date**—To view alerts from a specific period.
 - **Source Name**—To view the alerts from a specific system.
 - **Category** and **Subcategory**—To view alerts of specific category.
 - **Message**—To view alerts containing a specific word in the message column.

Selections that are made in the filters are applied at real time.

4. To reset the filters, click **Clear All Filters**.

Acknowledging alert logs

You can acknowledge alert logs that are not already acknowledged. Acknowledging an alert prevents storing the same event in the system. For example, if a device is noisy and is generating the same event multiple times, you can ignore further recording of the alert by acknowledging the events that are received from the device. And, no events of the same type are recorded further.

To acknowledge alert logs:

On the **Alert Log** page, select the alert logs that you want to acknowledge and click **Acknowledge**.

A check mark is displayed in the **Acknowledge** column for the selected alert logs.

Unacknowledging alert logs

You can unacknowledge alert logs that are acknowledged. Unacknowledging an alert implies that all events from any device are recorded even when the same event recurs frequently. By default, all alerts are unacknowledged.

To unacknowledge alert logs:

On the **Alert Log** page, select the alert log that you want to unacknowledge and click **Unacknowledge**.

The check mark that is displayed in the **Acknowledge** column for the selected alert logs is cleared, indicating that the selected alert logs are unacknowledged.

Ignoring alert logs

You can ignore alert logs when you do not want to record an alert. No actions are initiated for any events occurring in the device with which the alert is associated. Alert policies for the selected device contain details of the events that must be ignored.

To ignore alert logs:

On the **Alert Log** page, select the alert logs that you want to ignore and click **Ignore**.

A message is displayed indicating that an alert policy is created to ignore alert logs of the type you selected. The ignore policy is created from the device or multiple devices where the alert log is generated.

Exporting alert logs

You can export alert logs in `.csv` format to a network share or local drive on your system.

To export alert logs:

On the **Alert Log** page, select the alert logs that you want to export and click **Export > Export Selected**.

You can export all alert logs by clicking **Export > Export All**.

The alert logs are exported in `.csv` format.

Deleting alert logs

You can delete one or multiple alert logs.

To delete alert logs:

On the **Alert Log** page, select the alert logs that you want to delete and click **Delete**.

A message is displayed prompting to you confirm the action.

Alert policies

The alert policies feature enables you to view critical alerts and perform specific tasks. To view the list of alert policies, click **Alerts > Alert Policies**. The alert policy details include name and description of the alert policy, status of the alert policy, email ID of the administrator, and syslog.

You can perform the following tasks on the **Alert Policies** page:

- Create alert policies
- Edit alert policies
- Enable alert policies

- Disable alert policies
- Delete alert policies

OME—Modular also offers pre-defined alert policies for monitoring the systems, after the alert destinations are configured.

Creating alert policies

To create an alert policy:

1. From the menu bar, click **Alerts > Alert Policies > Create**.
The **Create Alert Policy** wizard is displayed.
2. Enter the name and description for the alert policy.
3. Select **Enable Policy** to activate the alert policy and click **Next**.
The **Category** tab is displayed.
4. Select all alert categories, or select the required option and click **Next**. The available categories are:
 - Application
 - Chassis
 - iDRAC
 - Network IOMs
 - Storage IOMs

You can expand each category to view and select the subcategories.

The **Devices** tab is displayed.

5. Select the required devices or device groups and click **Next**.
The **Date and Time** tab is displayed.
6. Select the date, time, and days on which the alerts must be generated and click **Next**.
The **Severity** tab is displayed.
7. Select the severity level and click **Next**.
The available options are:

- All
- Unknown
- Info
- Normal
- Warning
- Critical

The **Actions** tab is displayed.

8. Select the alert action and click **Next**. The available options are:
 - **Email (Enable)**—Click **Enable** to view the **Email Configuration** window where you can configure the email settings for the alert.
 - **SNMP Trap Forwarding (Enable)**—Click **Enable** to view the **SNMP Configuration** window where you can configure the SNMP settings for the alert.
 - **Syslog (Enable)**—Click **Enable** to view the **Syslog Configuration** window where you can configure the system log settings for the alert.
 - **Ignore**

You can view the alert policy attributes in the **Summary** tab.

Enabling alert policies

You can enable alert policies that are disabled. You can enable more than one alert policy at a time.

To enable alert policies:

On the **Alert Policies** page, select the alerts that you want to enable and click **Enable**.
A confirmation message is displayed.

Editing alert policies

You can edit alert policies.

To edit alert policies:

On the **Alert Policies** page, select the alerts that you want to edit and click **Edit**.
A confirmation message is displayed.

Disabling alert policies

You can disable alert policies that are enabled. You can disable more than one alert policy at a time.

To disable alert policies:

On the **Alert Policies** page, select the alerts that you want to disable and click **Disable**.
A confirmation message is displayed.

Deleting alert policies

You can delete alert policies that are enabled. You can delete more than one alert policy at a time.

To delete alert policies:

1. On the **Alert Policies** page, select the alerts that you want to delete and click **Delete**.
A message is displayed prompting you to confirm the action.
2. Click **Yes** to proceed.

Alert definitions

You can view description of the alert logs generated for events that associated with the chassis, and devices and components in the chassis, on the **Alerts Definition** page. The alert information that is displayed is as follows:

- Severity of the alert
- Message ID of the alert
- Alert message
- Category of the alert
- Subcategory of the alert

You can sort the list of alerts based on the **Advanced Filters**:

- **Message ID Contains**
- **Message Contains**
- **Category**
- **Subcategory**
- **Severity**

You can also select an alert to view the details on the right side of the **Alerts Definition** page. The details are—detailed description, recommended action, event source information, and criticality.

Filtering alert definitions

To filter alert definitions:

1. On OME–Modular web interface, navigate to **Alerts > Alert Definitions**.
2. Click **Advanced Filters**.
3. Select or update the following based on your requirement:
 - **Message Contains**—To view alerts containing a specific word in the message column.
 - **Message**—To view alerts containing a specific numeric or alphanumeric character.
 - **Category** and **Subcategory**—To view alerts of specific category.
 - **Severity**—To view all alerts with specific severity level.

Selections that are made in the filters are applied at real time.

4. To reset the filters, click **Clear All Filters**.

Monitoring audit logs

The audit log feature in OME–Modular enables you to monitor log entries related to:

- Log in attempts
- Appliance setup
- Chassis configuration change using RESTful API
- Change in alert filter configuration

On the **Audit Log** page, you can perform the following tasks:

- Sort the audit logs using the Advanced Filter.
- Export all the audit logs in `.csv` format to a network share or local drive on your system.

Quick Deploy audit logs are recorded as an overall operation, whenever they are created or updated. The quick deploy audit log details are similar to details of any other job that is created or updated in the system.

To view the **Audit Log** page:

From the menu bar, click **Monitor** > **Audit Logs**.

The **Audit Log** page is displayed.

Topics:

- [Filtering audit logs](#)
- [Exporting audit logs](#)
- [Monitoring jobs](#)

Filtering audit logs

To filter audit logs:

1. On the **Audit Logs** page, expand **Advanced Filters**.
2. Select or update the following based on your requirement:
 - **Severity**—To view audit logs of **Info**, **Warning**, **Critical**, or **All** severity levels.
 - **Start Time** and **End Time**—To view audit logs of a specific period.
 - **User**—To view audit logs from a specific user.
 - **Source Address**—To view audit logs from a specific system.
 - **Category**—To view audit logs of audit or configuration type.
 - **Description**—To view audit logs containing a specific word in the **Description** column.
 - **Message ID**—To view audit log containing a specific number or character

Selections made in the filters are applied at real time. To reset the filters click **Clear All Filters**.

Exporting audit logs

You can export selected or all audit logs in a `.csv` format to a local drive on your system or a network share.

To export audit logs:

1. On the **Audit Logs** page, select the audit logs that you want to export.
2. Click **Export**, and select **Export Selected**.
Else, you can click **Export** > **Export All**, to export all the audit logs.

Monitoring jobs

You can view the status of and details of jobs that are initiated in the chassis and its subcomponents, on the **Jobs** page. The jobs include firmware update and inventory refresh for devices.

To view the **Jobs** page, from the menu bar, click **Monitor > Jobs**.

You can perform the following tasks on the **Jobs** page:

- Filter jobs using **Advanced Filter**
- View a summary of the job.
- Run jobs
- Stop jobs
- Enable jobs
- Disable jobs
- Delete jobs

The job status is "Completed with errors", when one or more sub-tasks fail the request and the status is set to "Warning". If all the sub-tasks fail, status is "Failed". If all the tasks are completed successfully, the status is displayed as "Completed".

NOTE: When the "Lockdown mode" is enabled on iDRAC, the Blink LED job status for iDRAC is displayed as "failed" on the OME-Modular Jobs page, even though the job is successful in iDRAC.

Filtering jobs

To filter jobs:

1. On the **Jobs** page, click **Advanced Filter**.
2. Select or update the following based on your requirement:
 - **Status**—To view jobs based on status. The available options are:
 - Scheduled
 - Queued
 - Starting
 - Running
 - Completed
 - Failed
 - New
 - Completed with errors
 - Aborted
 - Paused
 - Stopped
 - Canceled
 - **State**—To view jobs based on state. The available options are:
 - Enabled
 - Disabled
 - **Job Type**—To view jobs based on the type. The available options are:
 - Debug Logs
 - Settings Update
 - Software Rollback
 - Device Action
 - Restore
 - Device Config
 - Chassis Profile
 - Inventory
 - Update
 - MCM OffBoarding
 - Backup
 - Profile Update
 - Quick Deploy

- MCM OnBoarding
 - MCM Group
 - **Last Run Start Date** and **Last Run End Date**—To view jobs based on the last run period.
- Selections made in the filters are applied at real time. To reset the filters click **Clear All Filters**.

Viewing job details

The Fabric Manager on-boarding is initiated when a Fabric Manager failover occurs in the IOM cluster. When a new Fabric Manager is discovered, OME - Modular initiates the on-boarding process to reestablish communication with the IOM cluster. In certain scenarios, multiple switchovers may occur within a short timespan resulting in failure of the tasks that are already in-progress. Only the last task is completed successfully. Following are the scenarios when multiple switchovers could occur:

- MM reset
- MM upgrade or switchover
- Inter-chassis link online insertion removal
- MM online insertion removal
- IOM Master upgrade
- IOM Master reset
- Fab-D congestions—Reasons for the congestion include downloading huge files that cause the FAB-D to drop other traffic

The details of the assigned MAC addresses for the respective NIC partitions are displayed on the **Jobs Details** page, based on the configuration results from iDRAC.

To view the details of a job:

1. On the **Jobs** page, select the job of which you want to view the details.
A summary of the job is displayed on the right side of the **Jobs** page.
2. Click **View Details**.
The **Job Details** page is displayed.

The details including name, description, execution details, and the details of the system on which the job was run, are displayed.

On **Job Details** page, you can perform the following tasks:

- **Restart** the job
- **Export** details of the job in a `.csv` format to a local drive on your system or a network share

 **NOTE: The Restart option for the MCM onboarding task for adding a member chassis is disabled irrespective of the job status.**

Sometimes after a firmware update, `racreset` or management module failover, a message stating that the alerts could not be retrieved is displayed. The message that is displayed does not impact the functionality of OME-Modular.

Running jobs

If a job is running for over 24 hours, stop the job after analyzing the job details. Rerun the job, if required.

You can use the **Jobs** page to run jobs immediately.

To run jobs:

On the **Jobs** page, select the jobs that you want to run and click **Run Now**.
A message is displayed to confirm that the task has restarted.

Stopping jobs

You can stop jobs that are in progress.

To stop jobs:

On the **Jobs** page, select the ongoing jobs that you want to stop and click **Stop**.
A message is displayed prompting you to confirm the operation.

Enabling jobs

You can enable jobs that are disabled.

To enable jobs:

On the **Jobs** page, select the disabled jobs that you want to enable and click **Enable**.
A confirmation message is displayed and the state of the selected jobs changes to "Enabled".

Disabling jobs

You can disable jobs that are enabled.

To disable jobs:

On the **Jobs** page, select the enabled jobs that you want to disable and click **Disable**.
A confirmation message is displayed and the state of the selected jobs changes to "Disabled".

Deleting jobs

To delete jobs:

On the **Jobs** page, select the jobs that you want to delete and click **Delete**.
A message is displayed prompting you to confirm the operation.

Use case scenarios

Assigning backup to the MCM Lead

The backup lead chassis feature facilitates management of systems in the chassis group when the existing lead chassis fails. Managing a chassis group consists of the following tasks:

- Assign—Allows assigning a member of the chassis group as a backup to the existing lead chassis.
- Unassign—Allows selection of another chassis in the group to replace the existing backup chassis.
- Promote—Allows the backup chassis to takeover as the lead chassis when the existing lead chassis fails.
- Retire—Allows the backup to takeover as the lead chassis when the existing lead chassis must be retired.

For more information, see [Chassis groups](#).

Lifecycle of backup

The life cycle of the backup feature consists of the following stages:

1. Stage 1—Creating a chassis group with backup lead.
2. Stage 2—Monitoring the health of the lead and backup.
3. Stage 3—Replacing the primary lead chassis with backup lead or retiring the lead chassis.

Creating chassis group with backup lead

To create a chassis group and assign a backup to the lead chassis, perform the following steps:

1. Rack and stack the chassis.
2. Wire multiple chassis in the rack. For more information, see [Wiring chassis](#) and [Pre-requisites for creating a distributed group](#).
3. Create a chassis group and add members to the group. For more information, see [Chassis groups](#).

Configuring a virtual IP is optional. The virtual IP enables a secondary IP on the lead that sticks with the lead. If the backup takes over as the new lead, then the secondary IP automatically moves to the new lead.

4. Configure the group from the lead chassis.

If there are any settings and configurations on the member chassis that could conflict with lead, clear those configurations before the lead pushes its configuration across the group. Do the following, if required:

- a. Configure chassis settings.
- b. Update firmware.
- c. Configure firmware baselines.
- d. Configure alert policies.
- e. Configure templates and identity pools, and deploy to devices or slots.
- f. Configure other settings.

5. Assign one of the members of the chassis group as the backup lead.

The initial configuration data synchronization from the lead chassis to backup chassis continues even after the assign job is completed. Both the lead and backup chassis report the health of the backup chassis.

Initially, the backup health status is displayed as "Critical" while the configuration data is being synchronized before changing to "OK". Wait for the backup health to transition to "OK" before proceeding. If the backup health continues to report "Critical" or "Warning" even after 30 minutes of the assign task, it is an indication that there are persistent communication issues. Unassign the backup and repeat the Step 5 to choose another member as the new backup. Also, Dell EMC recommends that you create an alert policy on lead to take notification actions through email, SNMP trap, system log, for backup health alerts. Backup health alerts are part of the chassis configuration and system health category.

6. Configure the member chassis that is designated as the backup.

It is mandatory for the backup chassis to have its own management network IP. The IP enables the backup to forward backup health alerts.

Create an alert policy on the backup to take notification actions (email, SNMP trap, system log) for backup health alerts. Backup health alerts are part of Chassis (Configuration, System Health) category. The backup chassis raises warning or critical alerts when it detects that the backup synchronization status is bad because of communication or other irrecoverable errors.

Monitoring the MCM group

1. Complete all the configuration tasks before assigning the backup lead. However, if you have to modify the configuration after assigning the backup, the changes are automatically copied to the backup. The process of copying the changes to the backup may take up to 90 minutes, based on the configuration change.
2. The backup synchronization status of the lead and backup lead chassis is available at the following GUI locations:
 - a. On the lead chassis:
 - **Home** page—**Backup Sync** status under the member (backup)
 - Lead **Overview** page—Redundancy and backup synchronization status under **Group Information**
 - b. On the backup chassis:
 - **Home > Overview** page—**Backup Sync** status under the **Group Information**.
3. Interpreting the backup health:
 - If backup sync is healthy, the status is displayed as "Ok" and no further actions are needed.
 - If backup sync is not healthy, the status is displayed as "Warning" or "Critical". The "Warning" indicates a momentary synchronization problem that is resolved automatically. The "Critical" status indicates a permanent problem and requires user action.
 - When the backup sync status changes to "Warning" or "Critical", the associated alerts are generated under alert categories Chassis (Configuration, System Health). These alerts are logged to the **Home > Hardware Logs** and **Alerts > Alert Log**. The alerts are also shown as faults under the **Home > Chassis Subsystems** (top right-hand corner) under the MM subsystem. If an alert policy is configured, the actions are taken as configured in the policy.
4. Required user actions when Backup health is "Warning" or "Critical":
 - Warning—A momentary status and must transition to "Ok" or "Critical". But if the status continues to report "Warning" for more than 90 minutes, Dell EMC recommends that you assign a new backup.
 - Critical—A permanent status indicative of issues with the backup or lead. Identify the underlying issues and take appropriate actions as described below:
 - Health is critical because of alert CDEV4006: The lead or member chassis has drifted its firmware version causing a lead/backup incompatibility. It is recommended that the firmware of the lead or member chassis is brought back to the same version (1.10.00 or later).
 - Health is critical because of alert CDEV4007: one of the several underlying issues contributes to this status, see the following flow chart to determine the cause and take the recommended action.




Figure 2. Network and power outage—flowchart

The alert, CDEV4007, is related to network or power issues that can be classified as:

- **Intermittent/recoverable issues**—Momentary power or network outages. The administrator can identify these types of failures and perform recovery actions locally or remotely. Do not promote the backup lead. Allow the lead chassis to recover connectivity automatically or the administrator fixes the power or network issues.
- **Partial failure**—Both management modules fail or malfunction. But the remaining chassis components are working. Promote the backup lead as the lead chassis to regain group management function through the new lead. For more information about promoting the backup and restoring the failed lead chassis to production state, see the section, [Disaster recovery of lead chassis](#).
- **Complete failure**—Catastrophic failures. All the chassis components including the management modules are broken or nonresponsive. Promote the backup lead as the lead chassis to regain group management function through the new lead. For information about promoting the backup lead and clearing references to the failed lead chassis, see the section, [Disaster recovery of lead chassis](#).

Scenarios when backup lead can take over as lead chassis

This section describes the situations in which a backup lead can take over as the lead chassis of the chassis group.

Disaster recovery of lead chassis

Catastrophic failures such as power loss, network loss, and failure of both MMs can result in the lead chassis being inaccessible or unavailable. In such cases, you can promote the backup to take over from the failed lead chassis for continued management of systems.

NOTE: Promoting the backup lead as the new lead restores the group management function for the member chassis that are not exposed to the failures. However, there are limitations on the extent of the functionality can be restored on the failed lead chassis. The restoration is based on the severity of the failures in the failed lead chassis.

Remember the following while recovering the lead chassis:

1. Before running the "promote" task on the backup lead chassis:
 - a. The "promote" task is a disruptive operation and must be used only when there are no means to recover the inaccessible lead chassis. In partial failures of the lead chassis, for example; if only the management modules are nonresponsive, but the computes

are working, running the promote task disrupts workloads that are still running on the lead chassis computes. For information about relocating working components that is, computes and network switches from the failed lead, see, the list item 3.c, "Steps that are required to restore the failed lead before putting it into production."

- b. After determining that the lead chassis has failed and is inaccessible, you must remotely shut down power to the lead chassis or physically remove the chassis from the stack before running the "promote" task on the backup. If lead chassis not turned off or removed from the stack before the promote task, the failed or partially failed lead chassis may revive after promoting the backup and cause situations of multiple leads. Multiple leads can create confusion and interference in managing the group.
2. Running the "promote" task on the backup lead chassis:
 - a. If the lead chassis is up and running, the backup chassis web interface blocks the "promote" task. Ensure that the lead has failed and is inaccessible before initiating the promote task on backup. The backup may erroneously block the "promote" when the lead is accessible through the private network, but it may not be reachable on the public user management network. In such cases, OME-Modular RESTful API can be used to run the promote task forcefully. For more information, see the RESTful API guide.
 - b. A job is created after the "promote" operation is started. The job may be completed in 10-45 minutes, based on the number of chassis in the group and amount of configuration that has to be restored.
 - c. If the lead chassis is configured to forward alerts to external destinations (email, trap, system log), any alerts that components in the group generate while the lead is down, are available only locally in their respective hardware or alert logs. During the lead outage, the leads cannot be forwarded to configured external destinations. The outage is the period between lead failure and successful promotion of backup.
 3. Expected behavior after the "promote" task:

- a. The backup chassis becomes the lead and all the member chassis are accessible as they were on the earlier lead chassis. After the "promote" task, references to the old lead chassis exist as a member of the same group. The references are created to prevent any disruption to the working computes in the old lead in a lead chassis MM failure situation.

The "promote" task rediscovers all the members in the group and if any member chassis is inaccessible then, the chassis is still listed in the lead home page with a broken connection and available repair options. You can use the repair option to add the member chassis again or remove the chassis from the group.

- b. All firmware baselines or catalogs, alert policies, templates or identity-pools, and fabrics settings are restored as they were on the failed lead chassis. However, following are some exceptions and limitations:
 - i. Any recent configuration changes on the failed lead within the 90 minutes window that is needed for copying to the backup, those configurations may not be copied completely to the backup and are not restored completely after the "promote" task.
 - ii. The in-progress and partially copied jobs that are associated with templates/identity-pools continue to run. You can perform one of the following tasks:
 - i. Stop the running job.
 - ii. Reclaim any identity-pool assignments.
 - iii. Restart the job to redeploy the template.
 - iii. Any template that is attached to an occupied slot through the lead before the backup takes over as the new lead, is not deployed on the existing sled when it is removed or reinserted. For the deployment to work, the administrator must detach the template from the slot, reattach the template to the slot, and remove or reinsert the existing sled. Or, insert a new sled.
 - iv. Any firmware catalogs that are created with automatic update catalog on a schedule are restored as manual updates. Edit the catalog and provide automatic update method with update frequency.
 - v. Alert Policies, with stale or no references to devices on the old lead, are not restored on the new lead.
- c. Steps that are required to restore the failed lead before putting it into production:
 - i. On the new lead, turn off the chassis remotely before performing the "promote" task on the backup. If the chassis not turned off, the partially failed lead may come online and cause a situation of multiple leads. There is limited support in automatic detection and recovery of this situation. If the earlier lead comes online and automatic recovery is possible, the earlier lead is forced to join the group as a member.
 - ii. On the new lead, remove the earlier lead chassis from the group to remove references to it.
 - iii. On the old lead, gain physical access to the failed lead chassis as soon as possible and unstack it from the group. If there were any templates with identity-pool assignments that are deployed to any computes on the old lead, then reclaim the identity-pool assignments from the computes. Reclaiming the identity pool assignments is required to prevent any network identity collision when the old chassis is put back into production.
 - iv. Do not delete fabrics from the old lead chassis as deleting the fabrics can lead to network loss once the old lead is added back to the network.
 - v. On the old lead, run a force "reset configuration" using the following REST API payload:

URI: /api/ApplicationService/Actions/ApplicationService.ResetApplication

Method: POST

Payload: {"ResetType": "RESET_ALL", "ForceReset": true}

- d. Relocate the working components of the old lead to other chassis in the group:

- i. Relocate network switches from the old lead to the new lead or member chassis in the group to restore the health of the fabrics.
- ii. Relocate computes from the old lead to the new lead or member chassis in the group. New templates or identities must be deployed on the computes before resuming workloads, which they were running on the old lead chassis.

Retiring lead chassis

The "retire" option enables a backup chassis to takeover as the lead of a chassis group when the lead chassis is running for a long time and must be removed from the production environment temporarily or permanently. The lead chassis can gracefully detach from the group. The "retire" option also facilitates the lead to retire from the lead role but remain a member of the group.

1. Run "retire" task from the lead chassis:
 - a. A job is created when the "retire" task starts. The job may be completed in 10-45 minutes based on the number of chassis in the group and amount of configuration to be restored.
 - b. If the lead chassis is configured to forward alerts to external destinations (email, trap, system log), any alerts that the components in the group generate are only available locally in their respective hardware or alert logs when the retire task and the backup taking over the lead chassis is in progress. After the "retire" task is complete and before the backup is promoted, there is an outage in group management. The outage includes forwarding of alerts to configured external destinations.
2. Expected behavior of backup on completion of the "retire" task:
 - a. The backup chassis becomes the new lead and all the member chassis are accessible as they were on the retired lead chassis. The new lead chassis rediscovers all the members in the group and if any member chassis is inaccessible then, the members are still listed on the **Home** page of the lead chassis with broken connection and available repair options. Use the repair option to readd or remove the member chassis from the group.
 - b. All firmware baselines or catalogs, alert policies, templates or identity-pools, and fabrics settings are restored, as they were on the retired lead chassis.
3. Expected behavior of old lead chassis on completion of the "retire" task:
 - a. If the old lead is chosen to retire as a stand-alone chassis, it continues to carry the templates/identity-pools configuration. Perform the following steps to clear configuration to avoid conflicts with the new lead.
 - i. Unstack the earlier lead from the group.
 - ii. Reclaim any identity-pool IO identities that are deployed to computes on the old lead.
 - iii. Do not delete fabrics from the old lead chassis as deleting the fabrics can lead to network loss once the old lead is added back to the network.
 - iv. Run a force "reset configuration" using the following REST API payload:

URI: `/api/ApplicationService/Actions/ApplicationService.ResetApplication`

Method: `POST`

Payload: `{"ResetType": "RESET_ALL", "ForceReset": true}`
 - b. If the old lead is retired as a member of the current group, it no longer carries the identity-pools configuration. However, it contains the templates configuration. To avoid conflicts with the new lead, clear the templates configuration using **Configuration > Deploy > Delete**.

Troubleshooting

This section describes the tasks for troubleshooting and resolving issues using the OME–Modular user interface.

- Firmware update is failing
- Storage assignment is failing
- Management role of IOMs is downgraded
- IOM health is downgraded
- Drives on compute sled are not visible
- Storage sleds cannot be applied to IOMs
- Drives in OpenManage are not visible
- iDRAC drive information does not match OpenManage drive information
- The assignment mode of storage sled is unknown

Topics:

- [Storage](#)
- [Unable to access OME-Modular using Chassis Direct](#)
- [Troubleshooting lead chassis failure](#)

Storage

This section describes the issues that are related to storage sleds and steps to resolve the issues.

Firmware update is failing

1. Firmware update may fail if one or more subcomponents fail to flash during the firmware update process.
2. If an IOM is downgraded owing to a chassis mismatch or faulty subcomponent, the firmware activation fails.

Storage assignment is failing

A storage assignment fails if:

1. The IOMs are currently downgraded.
2. There is only one IOM present.
3. Only one hot-swappable Expander is present on a storage sled.

SAS IOM status is downgraded

Both SAS IOMs are degraded if a:

1. Peer SAS IOM is detected but cannot be communicated with.
2. Firmware Mismatch is detected.
3. Chassis Mismatch is detected.

SAS IOM health is downgraded

The SAS IOM health is downgraded if:

1. One or more subcomponents are faulty.
2. A non-SAS IOM is detected.
3. An inconsistency is detected in the subcomponent firmware.

Drives on compute sled are not visible

1. If the compute sled is configured with a PERC controller and the drives have been reseated or moved, they are rediscovered as "Foreign".
2. If the drives are removed from the storage sled, they cannot be discovered.
3. If a storage sled is replaced, the storage configuration of the earlier sled cannot be applied to the replaced sled.

Storage configuration cannot be applied to SAS IOMs

1. If a storage sled is replaced, the storage configuration of the earlier sled cannot be applied to the replaced sled.
2. If a firmware mismatch is detected on the boot of the SAS IOM, the storage configuration is not applied.
3. If a chassis mismatch is detected on the boot of the SAS IOM, the storage configuration is not applied.
4. If the storage sled cannot be communicated with or has an Expander fault, the SAS IOM cannot apply the respective storage configuration.

Drives in OpenManage are not visible

1. The storage sled may have experienced an Expander failure which blocks the drives from being inventoried.
2. To view the drives, refresh the storage sled inventory.

iDRAC and OpenManage drive information do not match

The drive information of iDRAC and OpenManage may not match owing to the mechanisms that iDRAC and the SAS IOM used to fetch and detect the storage details for storage sleds.

The assignment mode of storage sled is unknown

1. If the IOM management role is currently downgraded, then the storage sled assignment mode may not be read.
2. You may have to refresh the **Storage** sled inventory page.
3. If the storage sled health is non-optimal the assignment mode may be downgraded.

Unable to access OME-Modular using Chassis Direct

On systems running Linux operating systems, you may be unable to access `ome-m.local` from your web browser using Chassis Direct. The inaccessibility could be due to missing IP address on the USB network link on the system. To fix this issue, perform one of the following steps while the USB cable is connected to the system and the chassis.

- On the system, go to the **Settings > Network** and enable **USB Ethernet**.
- On the top-right corner of the screen and click **Connect**.

Troubleshooting lead chassis failure

When a lead chassis is in phase of coming online after it has failed, the transition must be detected automatically. If you have promoted the backup lead chassis as the new lead chassis, ensure that the earlier lead chassis transitions properly before you put it back into the production environment.

Before putting the earlier lead chassis back into production, perform the following steps:

1. Disconnect the stacking cable.
2. Run the RESTful API to force reset to default.
The lead chassis becomes a stand-alone chassis.
3. Connect the stacking cable and add the stand-alone member to the same or different chassis group.

Recommended slot configurations for IOMs

The table below contains the recommended IOM slot configurations.

Table 13. Recommended IOM slot matrix

Slot A1	Slot A2	Slot B1	Slot B2
MX9116n	MX9116n	Empty	Empty
MX5108n	MX5108n	Empty	Empty
MX7116n	MX7116n	Empty	Empty
25G PTM	25G PTM	Empty	Empty
10GBT PTM	10GBT PTM	Empty	Empty
MX9116n	MX9116n	MX9116n	MX9116n
MX5108n	MX5108n	MX5108n	MX5108n
MX7116n	MX7116n	MX7116n	MX7116n
MX9116n	MX7116n	Empty	Empty
MX7116n	MX9116n	Empty	Empty
MX9116n	MX7116n	MX9116n	MX7116n
MX7116n	MX9116n	MX7116n	MX9116n
25G PTM	25G PTM	25G PTM	25G PTM
10GBT PTM	10GBT PTM	10GBT PTM	10GBT PTM

Topics:

- [Supported slot configurations for IOMs](#)

Supported slot configurations for IOMs

The table below contains the supported IOM slot configurations.

Table 14. Supported IOM slot matrix

Slot A1	Slot A2	Slot B1	Slot B2
MX9116n	Empty	Empty	Empty
MX5108n	Empty	Empty	Empty
MX7116n	Empty	Empty	Empty
25G PTM	Empty	Empty	Empty
10GBT PTM	Empty	Empty	Empty
MX9116n	Empty	MX9116n	Empty
MX5108n	Empty	MX5108n	Empty
MX7116n	Empty	MX7116n	Empty
25G PTM	Empty	25G PTM	Empty

Table 14. Supported IOM slot matrix(continued)

Slot A1	Slot A2	Slot B1	Slot B2
10GBT PTM	Empty	10GBT PTM	Empty
MX9116n	MX9116n	MX9116n	Empty
MX5108n	MX5108n	MX5108n	Empty
MX7116n	MX7116n	MX7116n	Empty
25G PTM	25G PTM	25G PTM	Empty
10GBT PTM	10GBT PTM	10GBT PTM	Empty
MX9116n	MX9116n	MX5108n	MX5108n
MX9116n	MX9116n	25G PTM	25G PTM
MX9116n	MX9116n	10GBT PTM	10GBT PTM
MX9116n	MX7116n	MX5108n	MX5108n
MX7116n	MX9116n	MX5108n	MX5108n
MX9116n	MX7116n	25G PTM	25G PTM
MX7116n	MX9116n	25G PTM	25G PTM
MX9116n	MX7116n	10GBT PTM	10GBT PTM
MX7116n	MX9116n	10GBT PTM	10GBT PTM
MX7116n	MX7116n	MX5108n	MX5108n
MX7116n	MX7116n	25G PTM	25G PTM
MX7116n	MX7116n	10GBT PTM	10GBT PTM
MX5108n	MX5108n	MX9116n	MX9116n
MX5108n	MX5108n	MX7116n	MX7116n
MX5108n	MX5108n	MX9116n	MX7116n
MX5108n	MX5108n	MX7116n	MX9116n
MX5108n	MX5108n	25G PTM	25G PTM
MX5108n	MX5108n	10GBT PTM	10GBT PTM
25G PTM	25G PTM	MX9116n	MX9116n
25G PTM	25G PTM	MX7116n	MX7116n
25G PTM	25G PTM	MX9116n	MX7116n
25G PTM	25G PTM	MX7116n	MX9116n
25G PTM*	25G PTM*	10GBT PTM*	10GBT PTM*
10GBT PTM	10GBT PTM	MX9116n	MX9116n
10GBT PTM	10GBT PTM	MX7116n	MX7116n
10GBT PTM	10GBT PTM	MX9116n	MX7116n
10GBT PTM	10GBT PTM	MX7116n	MX9116n
10GBT PTM*	10GBT PTM*	25G PTM*	25G PTM*

LEGEND:

*—Combining two types of Pass-Through Modules (PTMs) is supported.