

Dell EMC PowerFlex Rack

Using PowerFlex Manager

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.

Revision history.....	5
Chapter 1: Introduction.....	7
Chapter 2: Getting started.....	8
How PowerFlex Manager works with PowerFlex rack.....	8
Deployment modes.....	9
Managed mode.....	9
Lifecycle mode.....	9
Alerting mode.....	10
Sample templates.....	10
Networks.....	12
Network types.....	12
Define a network.....	12
Chapter 3: Monitoring system health.....	14
Use the PowerFlex Manager Dashboard.....	14
Monitor a service.....	16
Monitor resources.....	16
Resource health status.....	17
Resource compliance status.....	18
Resource operational state.....	18
Monitor PowerFlex Manager storage usage.....	18
Alerts.....	18
View alerts.....	19
Alert messages and recommended actions.....	19
Verify the alert connector status.....	20
Chapter 4: Managing system resources in services.....	21
Managing external changes.....	21
Service information.....	22
View service details.....	22
Edit service information.....	24
Service actions.....	25
Service mode options.....	26
Enter and exit service mode.....	26
Adding volumes.....	27
Add a volume to a PowerFlex hyperconverged node.....	27
Add a volume to a PowerFlex storage-only node.....	28
Add a volume to a PowerFlex compute-only node.....	28
Resize a volume.....	29
Add a network to a service	29
Reconfiguring MDM roles.....	30
Remove a service.....	30
Services deployed outside of PowerFlex Manager.....	32

Add a service deployed outside of PowerFlex Manager.....	33
Update resource inventory.....	34
Chapter 5: Discovering resources.....	35
Discover a resource.....	36
Chapter 6: Managing credentials.....	38
Update a credential in PowerFlex Manager.....	38
Update a credential for a PowerFlex node in a service.....	38
Create a credential.....	39
Edit a credential.....	40
Delete a credential.....	40
Updating passwords for system components.....	40
Chapter 7: Managing compliance.....	41
Compliance status.....	41
Update firmware and software.....	41
View a compliance report and perform remediation.....	42
Export a compliance report for all resources.....	43
Chapter 8: Updating PowerFlex Manager and components.....	44
Chapter 9: Expanding a system.....	45
Discover and configure a PowerFlex node for alerting.....	45
Ensure the availability of network IP addresses.....	46
Add PowerFlex nodes to a service.....	46
Chapter 10: Managing the PowerFlex Manager virtual appliance.....	48
Configure syslog forwarding.....	48
SNMP communication.....	48
Configure SNMP.....	49
Importing Active Directory users.....	50
Restarting the PowerFlex Manager virtual appliance.....	51
Chapter 11: Troubleshooting and support.....	52
Generate a troubleshooting bundle.....	52
Access the application logs.....	53
Exporting a configuration report for all resources and services.....	54

Revision history

Date	Document revision	Description of changes
January 2021	5.0	<p>Added the following new topics:</p> <ul style="list-style-type: none"> ● <i>Exporting a configuration report for all resources and services</i> ● <i>Restarting the PowerFlex Manager virtual appliance</i> ● <i>Updating passwords for system components</i> <p>Updated content for:</p> <ul style="list-style-type: none"> ● Monitoring resources ● Sample templates ● Network management ● System health and configuration checking ● Discovering resources ● Creating credentials with non-root user accounts ● Configuring SNMP ● Alerting for upgrade failures ● vCenter enhanced link mode support ● Fault set definitions
August 2020	4.0	<p>Added the following new topics:</p> <ul style="list-style-type: none"> ● <i>Sample templates</i> ● <i>Lifecycle mode</i> ● <i>Resizing a volume</i> ● <i>Discovering resources</i> <p>Updated content for:</p> <ul style="list-style-type: none"> ● Lifecycle mode ● Updating PowerFlex Manager and components ● PowerFlex protected maintenance mode ● PowerFlex replication ● Static route configurations ● Alerting ● Email alerts ● SNMP communication and configuring SNMP ● Expansion prerequisites ● System topology reporting ● Miscellaneous PowerFlex Manager 3.5 enhancements
March 2020	3	<p>Updated these topics:</p> <ul style="list-style-type: none"> ● <i>Expansion prerequisites</i> ● <i>Create a credential</i>
March 2020	2	<p>Added the following topics or sections:</p> <ul style="list-style-type: none"> ● <i>Verify the alert connector status</i> ● <i>Reconfiguring MDM roles</i> ● <i>Expansion prerequisites</i> ● <i>Managing the virtual appliance</i> <p>Updated these topics:</p> <ul style="list-style-type: none"> ● <i>Enter and exit service mode</i> ● <i>Monitor resources</i> ● <i>Services deployed outside of PowerFlex Manager</i> ● <i>Edit service information</i> ● <i>Add a service deployed outside of PowerFlex Manager</i>

Date	Document revision	Description of changes
		<ul style="list-style-type: none">• <i>View a compliance report</i>• <i>Define a network</i>• <i>Add a network to a service</i>• <i>Remove a service</i>• <i>Ensure the availability of network IP addresses</i>
January 2020	1	Initial release of the document

Introduction

This guide introduces PowerFlex Manager and the PowerFlex rack. It describes key tasks that are performed through PowerFlex Manager to manage your PowerFlex rack.

The guide assumes that Dell Technologies Services installed and configured PowerFlex Manager.

The target audience for this document is Dell EMC PowerFlex rack customers.

Dell EMC PowerFlex rack was previously known as Dell EMC VxFlex integrated rack. Similarly, Dell EMC PowerFlex Manager was previously known as Dell EMC VxFlex Manager, and Dell EMC PowerFlex was previously known as Dell EMC VxFlex OS. References in the documentation will be updated over time.

PowerFlex rack architecture is based on Dell EMC PowerEdge R640, R740xd, and R840 servers.

Dell EMC Networking S5200 series switches are now known as Dell EMC PowerSwitch S5200-ON series switches. Dell EMC Networking S4100 series switches are now known as Dell EMC PowerSwitch S4100-ON series switches. References in the documentation will be updated over time.

See the [Glossary](#) for terms, definitions, and acronyms.

You can also see the following locations for more information:

- *PowerFlex Manager online help* for additional documentation about using PowerFlex Manager. To access the PowerFlex Manager online help, click the **?** icon in the upper right of any screen in the UI.
- [Dell EMC PowerFlex Rack Administration Guide](#) to administer your PowerFlex rack.
- [Support](#) to search the knowledge base for FAQs, Tech Alerts, and Tutorials.

Getting started

PowerFlex Manager provides IT operations management for Dell EMC PowerFlex rack. It increases efficiency by reducing time-consuming manual tasks that are required to manage system operations. Use PowerFlex Manager to deploy and manage new and existing PowerFlex rack environments.

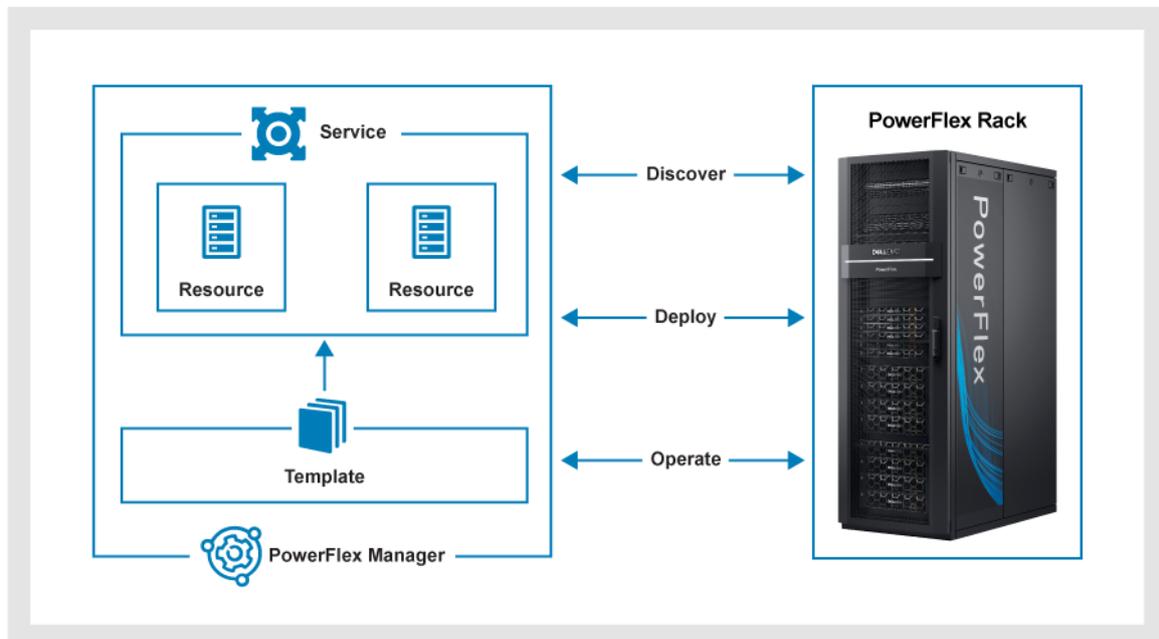
When you first log in to PowerFlex Manager, you see either the **Getting Started** page or the **Dashboard**.

The **Getting Started** page in PowerFlex Manager shows green marks for completed configurations steps. Only administrator users have access to this page.

If all configuration steps are complete, you land on the **Dashboard**, which provides system configuration details and communicates the health status for PowerFlex rack infrastructure elements and services.

How PowerFlex Manager works with PowerFlex rack

PowerFlex Manager discovers, deploys, and operates the PowerFlex rack by using resources, templates, and services.



The following table explains key PowerFlex Manager terminology.

Term	Description
Resource	The PowerFlex nodes, network switches, and virtual machine managers in the system. When PowerFlex Manager discovers resources, it identifies each component and how to talk to its management interface.
Template	Contains the configuration requirements that must be applied to a resource or group of resources during deployment. These requirements include firmware and software, operating system and/or hypervisor, and PowerFlex. A template represents the wanted state of your deployed configuration. PowerFlex Manager provides sample templates that represent the standard configuration options. You can clone a sample template and use it for your implementation. You can also define a custom template if necessary.

Term	Description
Service	<p>A group of resources that are managed within PowerFlex rack as a cluster. It is the state of your deployed configuration. PowerFlex Manager performs health monitoring, RCM compliance monitoring and remediation, and the ability to add or remove resources.</p> <p>A service is the finished outcome.</p>

Deployment modes

PowerFlex Manager is deployed in one of three modes: alerting, managed, or lifecycle.

You need a PowerFlex Manager license to deploy PowerFlex Manager in managed mode.

Mode	Description
Managed	Management and orchestration. PowerFlex Manager has deployed services or imported existing services for elements.
Lifecycle	Monitoring, service mode, and compliance upgrade. PowerFlex Manager puts a service in lifecycle mode if the configuration has limited support. All other service operations are blocked.
Alerting	Monitoring and alerting. PowerFlex Manager has not deployed services or imported existing services for elements.

Managed mode

PowerFlex Manager managed mode provides all the functionality of alerting mode.

It also provides the following functions:

- Deployment and configuration of PowerFlex nodes, switches, operating system provisioning, virtualization, PowerFlex storage, and CloudLink encryption
- Maintenance, including drive servicing, service mode, port views, and the ability to take over management of previously deployed PowerFlex rack environments
- Upgrade of PowerFlex node BIOS and firmware, PowerFlex, VMware ESXi, CloudLink Center and agents
- Expansion, including adding volumes, PowerFlex nodes, and switches to an existing environment

Lifecycle mode

Lifecycle mode allows lifecycle management of node components in an unsupported configuration. In lifecycle mode, a service supports health and compliance monitoring, service mode, and non-disruptive upgrade operations. All other service operations are blocked.

Lifecycle mode controls the operations that can be performed for configurations with limited support, including the following:

- No switch configuration
- No switch connectivity
- Invalid server inventory
- Missing network settings
- Unsupported server configurations
- Unsupported NIC or NIC teaming policies

PowerFlex Manager also puts a service in lifecycle mode if you select a minimal compliance version that includes firmware only for the service.

Alerting mode

A deployment is in alerting mode if there are no services. PowerFlex Manager alerting mode provides a number of functions.

- Supports discovery and inventory of system components (PowerFlex nodes, switches, VMware vCenter, CloudLink Center, and PowerFlex Gateway)
- Provides performance metrics for PowerFlex nodes, switches, and PowerFlex
- Monitors system component health status
- Indicates compliance status
- Sends alerts to Secure Remote Services and configures node iDRAC to send alerts to PowerFlex Manager

Sample templates

This topic lists the sample templates that are provided with PowerFlex Manager.

Except for the Compute Only - Windows template, you can use the sample templates with the Link Aggregation Control Protocol (LACP) bonding network architecture.

Template name	Features	Nodes	VMs	Data networks
Compute Only - ESXi - Partial Network	Standard ESXi compute-only nodes deployment with partial network automation. The partial network feature allows you to work with unsupported switches. If you choose to use the partial network feature, you give up the error handling and network automation features that are available with a full network configuration that includes supported switches. Depends on hyperconverged or storage-only nodes to be deployed first.	3	0	4
Compute Only - Linux - Partial Network	Standard Linux compute-only nodes deployment with partial network automation. The partial network feature allows you to work with unsupported switches. If you choose to use the partial network feature, you give up the error handling and network automation features that are available with a full network configuration that includes supported switches. Depends on hyperconverged or storage-only nodes to be deployed first.	3	0	4
Compute Only - ESXi	Standard ESXi compute-only nodes deployment. Depends on hyperconverged or storage-only nodes to be deployed first.	3	0	4
Compute Only - Linux	Linux-based compute-only nodes deployment. No volume mapping. Depends on hyperconverged or storage-only nodes to be deployed first.	4	0	4
Compute Only - Windows	Windows-based compute-only nodes deployment. No volume mapping. Depends on hyperconverged or storage-only nodes to be deployed first.	3	0	2

Template name	Features	Nodes	VMs	Data networks
Hyperconverged	Standard hyperconverged nodes deployment.	4	0	4
Hyperconverged - 100Gb	100 GB networking deployment. Nodes require Mellanox CX5 interface cards.	4	0	4
Hyperconverged - Compression	Compression-enabled deployment. Requires NVDIMM in nodes.	4	0	4
Hyperconverged - Encryption	CloudLink encryption deployment. Requires CloudLink Center in environment.	4	0	4
Hyperconverged - Leaf-Spine	Networking without virtual port channel deployment. This is the most common leaf-spine configuration.	4	0	4
Hyperconverged - Partial Network	Hyperconverged deployment with partial network automation. The partial network feature allows you to work with unsupported switches. If you choose to use the partial network feature, you give up the error handling and network automation features that are available with a full network configuration that includes supported switches.	4	0	4
Management - CloudLink Center	Two node CloudLink Center cluster deployment, for use in encryption templates.	0	2	NA
Management - PowerFlex Gateway	PowerFlex Gateway VM deployment, for use in all templates.	0	1	4
Management - Presentation Server	Presentation server VM deployment. The presentation server hosts the PowerFlex GUI.	0	1	NA
Storage Nodes	Standard storage-only nodes deployment.	4	0	4
Storage Nodes - 100Gb	100 GB networking deployment. Nodes require Mellanox CX5 interface cards.	4	0	4
Storage Nodes - Compression	Compression-enabled deployment. Requires NVDIMM in nodes.	4	0	4
Storage Nodes - Partial Network	Partial network deployment. The partial network feature allows you to work with unsupported switches. If you choose to use the partial network feature, you give up the error handling and network automation features that are available with a full network configuration that includes supported switches.	4	0	4
Storage Nodes - Replication	Replication-enabled deployment. Configures storage data replicator (SDR) and journal capacity.	4	0	4

Networks

Network types

You can manage various network types in PowerFlex Manager.

Network type	Use
General purpose LAN	Accesses network resources for basic networking activities.
Hypervisor management	Identifies the management network for a hypervisor or operating system that is deployed on a PowerFlex node.
Hypervisor migration	Manages the network that you want to use for live migration. Live migration enables you to move running virtual machines from one node of a failover cluster to a different node in the same cluster.
OS installation	Allows static or DHCP network for operating system imaging on PowerFlex nodes.
Hardware management	Out-of-band management of hardware infrastructure.
PowerFlex data	Data traffic between Storage Data Servers (SDS) and Storage Data Clients (SDC).
PowerFlex data (Storage Data Client traffic only)	Storage Data Client traffic only.
PowerFlex Data (Storage Data Server traffic only)	Used for Storage Data Server traffic only.
PowerFlex replication	Supports PowerFlex replication.
PowerFlex management	PowerFlex system management.

Define a network

Defining a network enables information to be populated in PowerFlex Manager, but it does not change your environment.

Prerequisites

Ensure that the following conditions are met before defining the network:

- A new VLAN is created on any switches that need access to that VLAN and is added to any management cluster server facing ports. The VLAN is then added to any northbound trunks to other switches that it must communicate with.
- PowerFlex Manager can communicate with the out-of-band management network.
- PowerFlex Manager can communicate with the operating system installation network in which the appliance is deployed.
- PowerFlex Manager can communicate with the hypervisor management network.
- The DHCP node is fully functional with the appropriate PXE settings to PXE boot images from PowerFlex Manager in your operating system installation network.

Steps

1. On the menu bar, click **Settings > Networks**.
2. Click **Define**.
3. In the **Name** field, enter the name of the network. Optionally, in the **Description** field, enter a description for the network.
4. From the **Network Type** drop-down list, select a network type.

For a hyperconverged architecture deployment with two data networks, you typically have two PowerFlex data networks that are defined with the PowerFlex data network type. The PowerFlex data network type supports both SDC and SDS communications. The PowerFlex data network type is used with hyperconverged nodes.

5. In the **VLAN ID** field, enter a VLAN ID between **1** and **4094**.

 **NOTE:** PowerFlex Manager uses the VLAN ID to configure I/O modules to enable network traffic to flow from the PowerFlex node to configured networks during deployment.

6. Optionally, select the **Configure Static IP Address Ranges** check box, and then do the following:

 **NOTE:** The **Configure Static IP Address Ranges** check box is not available for all network types. For example, you cannot configure a static IP address range for the operating system installation network type. Also, you cannot select or clear this check box to configure static IP address pools after a network is created.

- a. In the **Subnet** box, enter the IP address for the subnet. The subnet is used to support static routes for data and replication networks.
- b. In the **Subnet Mask** box, enter the subnet mask.
- c. In the **Gateway** box, enter the default gateway IP address for routing network traffic.
- d. Optionally, in the **Primary DNS** and **Secondary DNS** fields, enter the IP addresses of primary DNS and secondary DNS.
- e. Optionally, in the **DNS Suffix** field, enter the DNS suffix to append for hostname resolution.
- f. To add an IP address range, click **Add IP Address Range**. In the row, indicate the role in PowerFlex nodes for the IP address range and then specify a starting and ending IP address for the range. For the **Role**, select either:
 - **SDS or SDC:** Default; range is assigned to Storage Data Server (SDS) and Storage Data Client (SDC) roles.
 - **SDC:** Range is assigned to SDC role on PowerFlex hyperconverged nodes and PowerFlex compute-only nodes.
 - **SDS:** Range is assigned to SDS role on PowerFlex hyperconverged nodes and PowerFlex storage-only nodes.

Repeat this step to add IP address ranges based on the requirement. For example, you can use one range for SDC and another for SDS.

 **NOTE:** IP address ranges cannot overlap. For example, you cannot create an IP address range of 10.10.10.1–10.10.10.100 and another range of 10.10.10.50–10.10.10.150.

7. Click **Save**.

Monitoring system health

The PowerFlex Manager **Dashboard** dynamically shows the health status of a deployed system.

PowerFlex Manager receives alerts from discovered components and displays them on the **Alerts** page. It receives SNMP alerts directly from iDRAC and forwards them to Secure Remote Services. You can configure Dell EMC Networking OS10 switches to automatically send alerts to PowerFlex Manager.

Use the PowerFlex Manager Dashboard

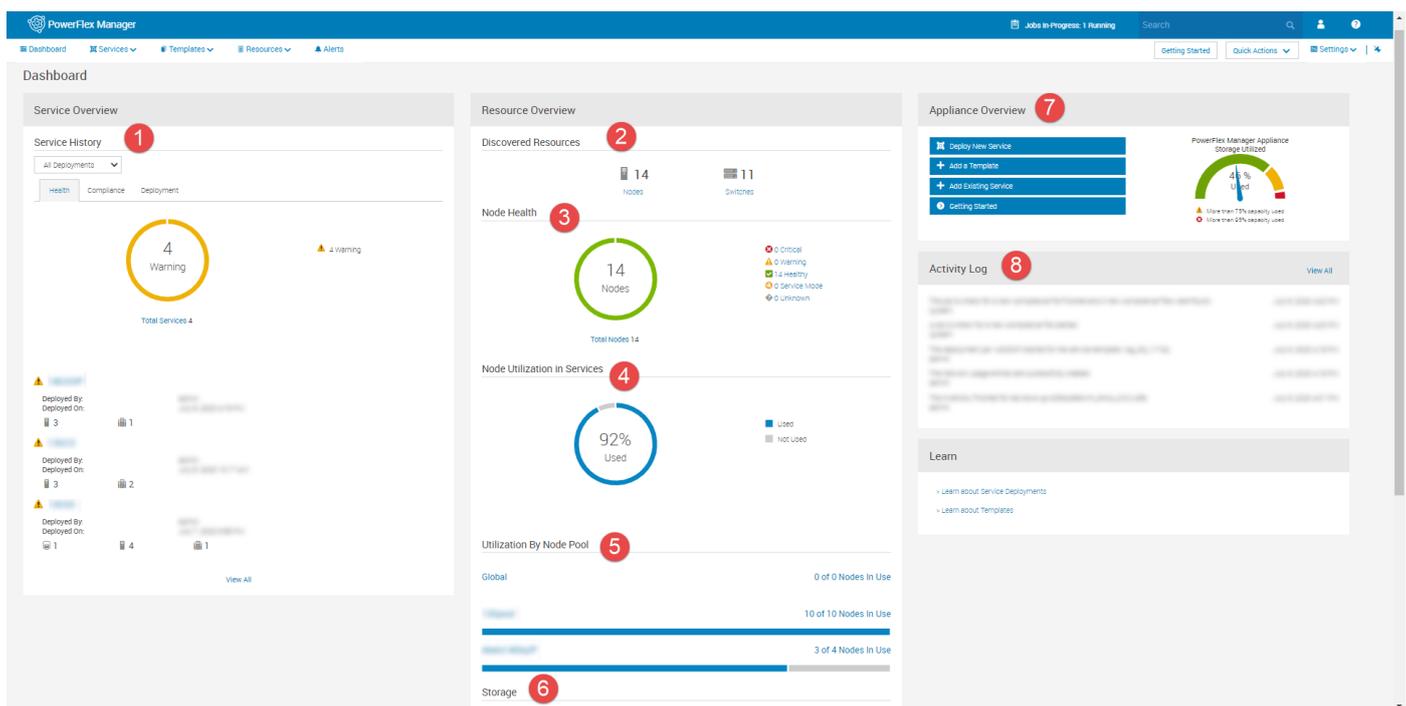
The Dashboard displays at-a-glance information about the system and its operational state. It provides the health status of the hardware and services that are deployed and running in the environment.

Use the Dashboard for a quick overview of the system, including:

- Health overview of the PowerFlex rack and individual PowerFlex nodes
- Usage of nodes, storage, and services
- Recently completed activities

Standard and operator users see only the details of services that they have created or for which they have permission.

Use the following figure to familiarize yourself with the Dashboard and its features.



Dashboard section	Description
1	<p>The Service History section displays a graphical representation of the status of deployed services. There are three tabs available:</p> <ul style="list-style-type: none"> • Health: Indicates by color whether the service is in a healthy (green band), warning (yellow band), or critical (red band) state. • Compliance: Indicates by color whether the service is compliant (green band) or noncompliant (yellow band) with the target compliance version.

Dashboard section	Description
	<ul style="list-style-type: none"> Deployment: Indicates by color whether the service deployed successfully (green band), is incomplete due to errors (red band), is being deployed (blue band), has no associated volumes (orange band), is in service mode (yellow band), or is in lifecycle mode (light green band). <p>To view a list of the services in a particular state, click the color band on the graphic.</p> <p>To view the status of a failed node component, hover the cursor on the image of the failed component in the service.</p>
2	<p>The Discovered Resources section shows the number of nodes and switches that are discovered in the system.</p> <p>To view detailed information about discovered resources, click the Nodes or Switches link.</p>
3	<p>The Node Health section graphically represents the node health through color:</p> <ul style="list-style-type: none"> A green band indicates the number of nodes that were successfully deployed. A yellow band indicates the number of nodes requiring corrective action but not affecting overall system health. It also indicates the number of nodes that have been placed in service mode. A red band indicates the number of nodes for which status is critical due to errors. A gray band indicates the number of nodes for which status is unknown. <p>To view a list of nodes in a particular state, click the color band on the graphic.</p>
4	<p>The Node Utilization in Services section displays a pie chart that shows:</p> <ul style="list-style-type: none"> The number of nodes in use (blue band) The percentage of nodes available for deployment (gray band) <p>To view the exact number, hover the cursor over each band.</p>
5	<p>The Utilization by Node Pool section displays the number of nodes that are used and available in a node pool. Each bar represents a node pool.</p>
6	<p>The Storage section displays a pie chart that shows:</p> <ul style="list-style-type: none"> The percentage of storage disk space in use (blue band) The percentage of disk space available (gray band) <p>To view the percentage, hover the cursor over each band.</p> <p>The Storage section also displays a count of storage components, including protection domains, volumes, data clients, and PowerFlex nodes.</p>
7	<p>The Appliance Overview section displays a gauge meter that shows the storage that is used by PowerFlex Manager:</p> <ul style="list-style-type: none"> A green status indicates that storage usage is below 75 percent. A yellow status indicates that storage usage is 75–95 percent. A red status indicates a critical storage situation, with storage usage above 95 percent. <p>This section also contains quick links to common PowerFlex Manager actions.</p>
8	<p>The Activity Log section lists the most recent user and system initiated activities.</p> <p>To view logs for all PowerFlex Manager application activities, click View All.</p>

Monitor a service

You can monitor the status of services from the **Dashboard** or from the **Services** page.

The following table describes the states that might appear for a service:

Icon	Service state	Description
	Healthy	The service was deployed successfully, and the resources are firmware compliant and healthy.
	Warning	One or more resources in the service require corrective action. This state does not affect the overall system health.
	Critical	Service deployment failed due to some issues.
	Pending	Service deployment is scheduled for a later time or date.
	In Progress	Service deployment is in progress.
	Canceled	Service deployment is canceled.
	Incomplete	The service is not fully functional because it has no associated volumes. Click Add Resources to add volumes.
	Service Mode	The service has been placed in service mode because one of the nodes within the service was put into service mode.
	Lifecycle Mode	The service supports health and compliance monitoring, service mode, and non-disruptive upgrades. All other service operations are blocked. Lifecycle mode controls the operations that can be performed for configurations that have limited support.
	Managed Mode	The service supports health and compliance monitoring, non-disruptive upgrades, automated resource addition, and automated resource replacement features.

On the **Services** page, you can view services in a particular state by selecting a state from the **Filter By** drop-down list.

Click a particular service to view more details, including cluster information, PowerFlex gateways, CloudLink details, storage volumes, and physical nodes. You can also view the overall service health, health of resources, resource compliance status, and deployment state.

Monitor resources

You can monitor the status of resources that PowerFlex Manager interacts with from the **Dashboard** or from the **Resources** page. Resources include nodes, network switches, virtual machine managers (for example, VMware vCenter), PowerFlex Gateway, and element managers (for example, CloudLink Center).

The **Resources** page displays information about all the resources that are discovered and managed in PowerFlex Manager, including the resource health status, compliance status, and operational state.

Click a resource to view detailed information about the resource and associated components. The information that displays depends on the resource type.

Resource type	Description
PowerFlex nodes	View performance details, including system usage, CPU usage, memory usage, and I/O usage. Performance usage values are updated every five minutes.
PowerFlex Gateway	View performance metrics, storage, and node details. You can export all metrics collected for a PowerFlex Gateway to a CSV file. The file contains aggregated values for hour, day, week, month, and year:

Resource type	Description
	<ul style="list-style-type: none"> Hourly: shows an entry for every five minutes for the last hour Daily: shows an entry for every hour for the last 24 hours Weekly: shows an entry for every day for the last seven days Monthly: shows an entry for every day for the last month Yearly: shows an entry for every month for the last 12 months
PowerFlex rack cabinet	View information about the cabinet, including system name and location, model number, IP address, cabinet temperature and humidity details.
Supported switches	View packet count by port for all ports or for a specific port. The packet count for all ports shows the average packet count for all the ports in that switch.

Resource health status

The following table describes the health states that might appear for a resource.

Icon	Health status	Description
	Healthy	The resource is working as expected.
	Warning	<p>The resource is in a state that requires corrective action, but does not affect the overall system health. For example, the firmware running on the resource is not at the required level or is not compliant.</p> <p>For supported switches, a warning health status might be because the SNMP community string is invalid or not set correctly. If a warning health status displays, PowerFlex Manager is unable to perform health monitoring for the switch.</p> <p> NOTE: If the service has VMware ESXi and is set to maintenance mode, a message displays in the View Details pane.</p>
	Critical	<p>An issue requiring immediate attention exists in one of the following hardware or software components:</p> <ul style="list-style-type: none"> Battery CPU Fans Power supply Storage devices Licensing <p>For supported switches, a critical health status might indicate that the power supply is not working properly or a CPU has overheated.</p> <p> NOTE: If the service is in a power off state, a message displays in the View Details pane.</p>
	Unknown	<p>The resource status is unknown.</p> <p>For example, for a supported switch, an unknown status might indicate that the switch has been turned off or is still being discovered.</p>
	Pending	The resource status is pending.
	Service Mode	The resource is in Service Mode .

Icon	Health status	Description
		 NOTE: If the resource fails to exit Service Mode , an informational message displays in the recent activity.

Resource compliance status

PowerFlex Manager assigns a firmware status to each resource.

Firmware status	Description
Compliant	No action is needed. The firmware running on the resource is compliant with the firmware version that is specified in the default compliance version.
Non-compliant	A firmware update is required. The firmware running on the resource is less than or greater than the firmware specified in the default compliance version.
Update required	A firmware update is required. The firmware running on the resource is less than the minimum firmware version recommended in the default compliance version.

Resource operational state

PowerFlex Manager tracks which resources it is managing and displays the operational state in the **Managed State** column on the **All Resources** tab.

Managed state	Description
Managed	PowerFlex Manager manages the firmware on the resource and can use the resource for deployments.
Unmanaged	PowerFlex Manager does not manage the resource. Unmanaged is the default state for all discovered nodes. To perform firmware updates or deployments on a discovered node, select the node and change the state to Managed . If PowerFlex Manager is configured for monitoring and alerting only, all resources are restricted to the Unmanaged resource state.
Reserved	PowerFlex Manager manages the firmware on the resource but cannot use the resource for deployments.

Monitor PowerFlex Manager storage usage

If PowerFlex Manager storage usage exceeds 75 percent, a banner displays at the top of the dashboard.

The banner indicates the current storage area of the partition with the least amount of space available (for example, /var, /tmp, /lib, or /home). The banner shows:

- A yellow warning status if storage usage is between 75 and 95 percent.
- A red critical status if storage usage is above 95 percent.

To free up disk space, you can delete compliance or operating system repository files from **Settings > Compliance and OS Repositories**.

You can hide the status message by selecting the hide option in the **Actions** menu.

Alerts

You can configure PowerFlex Manager to receive and display alerts from discovered PowerFlex rack components.

The alert connector is available through PowerFlex Manager. It sends alerts on the health of PowerFlex nodes securely through Secure Remote Services. Secure Remote Services routes alerts to the Dell EMC support queue for diagnosis and dispatch.

During node discovery, you can configure iDRAC nodes to automatically send alerts to PowerFlex Manager. PowerFlex Manager receives SNMP alerts directly from iDRAC and forwards them to Secure Remote Services.

If not done at discovery, you can configure iDRAC nodes to automatically send alerts to PowerFlex Manager by editing the alert connector settings and selecting the **Configure nodes for alert connector** option.

You can configure Dell EMC Networking OS10 switches to automatically send alerts to PowerFlex Manager by selecting the automatic configuration option during discovery.

You must manually configure CloudLink to send alerts to PowerFlex Manager.

You can configure the alert connector to send email alerts to one or more recipient email addresses you provide. You can receive these alerts in addition to, or instead of, Secure Remote Services alerts.

When the system is registered for alerting, a system topology report is automatically sent weekly to Secure Remote Services.

View alerts

You can configure PowerFlex Manager to receive alerts from discovered PowerFlex nodes. The alerts display on the **Alerts** page in PowerFlex Manager.

About this task

The **Alerts** page shows only node alerts. It does not include PowerFlex alerts. If you configure the alert connector to include email notifications, these email notifications include both node alerts and PowerFlex alerts.

Each generated alert displays the following fields:

Message ID	Acknowledged	Severity	Message	Category
Contains a prefix and number that is unique to each alert	Displays an icon when the alert is acknowledged Hover over the icon to see the user who acknowledged the alert	 Critical indicates that you should take immediate action to address the issues	Describes the event	Indicates the alert category: configuration, system health, updates, or miscellaneous
		 Unknown error indicates either the action that you should take to address the issue or a recommended work-around to reduce impact		
		 Warning indicates that you should schedule time to address the issue		
		 Informational indicates that the event requires no action		

For a list of alert messages and recommended response actions, see [Alert messages and recommended actions](#) on page 19.

PowerFlex Manager stores and displays up to 5 GB of SNMP alerts. Once this threshold is exceeded, PowerFlex Manager automatically purges the oldest data to free up space.

Steps

On the menu bar, click **Alerts**.

Alert messages and recommended actions

Use the code in the alert messages PowerFlex Manager generates to identify and correct issues. The alert message data contains information specific to the generated alert.

Message ID	Message	Severity	Recommended response action
VXOS0001	SDS has been in maintenance mode for more than 30 minutes.	Error (4)	Exit maintenance mode for the SDS in the PowerFlex cluster.

Message ID	Message	Severity	Recommended response action
	<SDS-ip> <PowerFlex Gateway>-<Service Tag>		
PFXM01	Appliance storage is within optimal parameters.	Info	No action is needed.
PFXM02	Storage usage is between 75 and 95 percent. It is recommended that unnecessary files be removed.	Warning	Remove unnecessary compliance files or operating system images. Consider disk expansion if additional components are added to the cluster.
PFXM03	Storage usage is above 95 percent and at a critical state. Remove any unnecessary files from the appliance.	Critical	Remove unnecessary compliance files or operating system images. Consider disk expansion if additional components are added to the cluster. Contact Dell EMC Support if space remains critical.
PFXM04	The operation cannot be completed because of an error during a firmware or software update.	Critical	<p>Retry the operation. If the problem persists, download the latest update package from your service provider and retry the operation again.</p> <p>To determine the reason for the error, look at the user interface logs or the console logs. PowerFlex Manager provides the name of the resource that failed, so look for logs that are related to that resource to identify the error.</p>
PFXM05	The operation cannot be completed because of an error during a PowerFlex Manager virtual appliance update.	Critical	<p>Check the network connection and access to the update package. Retry the operation. If the error persists, contact your service provider.</p> <p>To determine the reason for the error, look at the user interface logs or the console logs.</p>

Verify the alert connector status

You can verify the status of the alert connector to ensure that it is online. When the alert connector is online, alerts are being sent on system health through Secure Remote Services.

About this task

Status	Description
Online	The alert connector is online, and alerts are being sent.
Error	Contact Dell EMC Support. There is an issue and alerts are not being sent. For example, incorrect credentials are entered.
Missing	Contact Dell EMC Support. There is an issue and alerts are not being sent. For example, there is a communication issue between PowerFlex Manager and Secure Remote Services.

Steps

1. On the menu bar, click **Settings > Virtual Appliance Management**.
2. In the **Alert Connector** section, verify the status.

Managing system resources in services

You can view and interact with a service in PowerFlex Manager to perform compliance, health, and lifecycle management on the deployed infrastructure. PowerFlex Manager can discover PowerFlex rack environments that were deployed outside of PowerFlex Manager.

PowerFlex Manager supports the deployment of services that can contain any of the following types of PowerFlex nodes:

- PowerFlex hyperconverged node
- PowerFlex storage-only node
- PowerFlex compute-only node

Managing external changes

If you make manual changes outside of PowerFlex Manager, you must perform some steps within PowerFlex Manager. The steps ensure that the external changes are reflected within the user interface and the environment is kept in a healthy state.

The following table describes common tasks for managing external changes and what steps to take in PowerFlex Manager:

If you have...	Do this in PowerFlex Manager...
Manually replaced a node outside of PowerFlex Manager for a service in PowerFlex Manager	<ol style="list-style-type: none"> 1. Remove the node from the service. On the Services page, select the service. On the Details page, under Resource Actions, click Remove Resource. 2. Remove the node from the list of resources. On the Resources page, click the All Resources tab. From the list of resources, select the resource, and click Remove. 3. Discover the new node. On the Resources page, click Discover on the All Resources tab. 4. Update the service details. On the Services page, click Update Service Details.
Manually created additional VLANs on the VDS in vCenter outside of PowerFlex Manager	<ol style="list-style-type: none"> 1. Remove the service. On the Service Details page, click Remove Service under Service Actions. When you click Remove Service, do not select Delete Service because the Service Removal Type deletes the service and makes configuration changes to the nodes, switch ports, virtual machine managers, and PowerFlex. Instead, select Remove Service as the Service Removal Type. Then, to keep the nodes in the inventory, select Leave nodes in PowerFlex Manager inventory and set state to and choose Managed. For more information about removing a service, see <i>Removing a service</i> in the PowerFlex Manager online help. 2. Rediscover the service. On the Services page, click +Add Existing Service. <p>If you do not choose the correct options when you remove the service, you could delete the service and destroy it, or leave the servers in an unmanaged state and not be able to add the existing service.</p>
Manually created volumes outside of PowerFlex Manager	Click Update Service Details on the Services page for any service that requires the updated volumes.

If you have...	Do this in PowerFlex Manager...
Manually deleted volumes outside of PowerFlex Manager	Click Update Service Details on the Services page for any service that requires updated information about volumes deleted. PowerFlex Manager displays a message indicating that the volumes have been removed from the service.
Manually added nodes outside of PowerFlex Manager	Click Update Service Details on the Services page for any service that requires the updated nodes.
Manually removed nodes outside of PowerFlex Manager	Click Update Service Details on the Services page for any service that requires updated information about nodes deleted. Then, manually remove the resource on the Resources page by clicking the All Resources tab, selecting the resource, and clicking Remove .
Renamed objects such as ESXi host, volume, datastore, VDS, port group, data center, cluster, and so forth	Click Update Service Details on the Services page for any service that requires the updates.
Manually changed the IP address for a switch	<ol style="list-style-type: none"> 1. Remove the switch. On the Resources page, click the All Resources tab. From the list of resources, select the switch, and click Remove. 2. Rediscover the switch. On the Resources page, click the All Resources tab. Then, click Discover. 3. If the IP address is in use, remove the service and add it again. On the Services page, click +Add Existing Service.

Service information

You can view information about a service such as the overall service health, deployment details, and user permissions.

View	To find out
Overall Service Health	The aggregate state of health of the service.
Resource Health	The state of health of resources.
Compliance	Whether the resources are firmware or software-compliant when firmware or software updates are enabled on the service.
Deployment State	Whether the deployment of the service successfully completed.
Deployed By	The name of the user who deployed the service.
Deployed On	The date and time when the service was deployed.
Reference Template	The name of the reference template that is used in the service. NOTE: For existing services, the name displays as User Generated Template and not a template name from the inventory.
User Permissions	Whether user permissions are granted for one or more standard users to deploy the service.

View service details

You can view the state of a service at the component level. You can view the component topology and connections in a selected service template. You can view component logs, display the port view, place a component in service mode, or replace a disk.

About this task

If the service state is incomplete, a banner at the top of the page indicates that you must add volumes to make it fully functional. For more information about service states, see [Monitor a service](#) on page 16.

Steps

1. On the menu bar, click **Services**.
2. Under **All Services**, click the component whose details you want to view.
3. On the **Service Details**, click a component icon to view the following details associated with the component:
 - Management IP address
 - Resource health
 - Compliance
 - Deployment status
4. When you scroll down the page, the following information displays based on the resources in the service:

Section	Description
Clusters	View the following information about the clusters in the VMware vCenter environment: <ul style="list-style-type: none"> ● Health ● Data center name ● Cluster name ● Asset/Service tag ● Management IP address
PowerFlex Gateways	View the following information about the gateways that are part of the service: <ul style="list-style-type: none"> ● Health ● System name ● Primary MDM IP address ● Management IP address ● Virtual IP address ● Total storage pools <p>If the service has storage pools, PowerFlex Manager lists the number of storage pools, and the granularity setting and acceleration pool for each storage pool. If compression is disabled for a storage pool, the granularity is set to medium. If compression is enabled, the granularity can be set to fine or medium.</p> <p>The service deployment creates of a separate storage pool for each type of disk (SSD/NVMe or hard drive) that is found in the nodes. The deployment process adds the disks from the nodes to the appropriate storage pools based on the expected types for each pool.</p> <ul style="list-style-type: none"> ● Protection domain <p>In PowerFlex, a protection domain is a logical entity that contains a group of Storage Data Servers (SDS) that provide backup to each other. Each SDS belongs to one (and only one) protection domain. Each protection domain is a unique set of SDSs.</p>
CloudLink	View information about the CloudLink Center participating in the service, and the machine group and keystore being used for encryption.
Storage	View details about storage volumes added for the service. <p>The deployment of a PowerFlex hyperconverged node or PowerFlex storage-only node results in the creation of separate volumes within the storage pools. The volumes are mapped to storage pools according to the type of storage (SSD/NVMe or hard drive) expected for the storage pools.</p> <p>PowerFlex Manager includes the following information:</p> <ul style="list-style-type: none"> ● Health ● Volume ● Size ● Type ● Compression ● Storage pool ● Datastore

Section	Description
Physical Nodes	<p>View the following information about the nodes that are part of the service:</p> <ul style="list-style-type: none"> ● Health ● Asset/Service tag ● iDRAC management IP address ● Hostname ● PowerFlex mode <p>The mode for each node is one of the following:</p> <ul style="list-style-type: none"> ○ Hyperconverged includes both SDS and SDC components. ○ Storage Only includes only the SDS component. ○ Compute Only includes only the SDC component. <ul style="list-style-type: none"> ● Associated IP addresses ● MDM role <p>The MDM role is the metadata manager role. The MDM role applies only to those nodes that part of a PowerFlex cluster. The MDM role is one of the following:</p> <ul style="list-style-type: none"> ○ Primary: The MDM in the cluster that controls the SDSs and SDCs. The primary MDM contains and updates the MDM repository, the database that stores the SDS configuration, and how data is distributed between the SDSs. This repository is constantly replicated to the secondary MDMs, so they can take over with no delay. <p>Every PowerFlex cluster has one primary MDM.</p> <ul style="list-style-type: none"> ○ Secondary: An MDM in the cluster that is ready to take over the primary MDM role if necessary. ○ Tiebreaker: An MDM whose only role is to help determine which MDM is the primary. ○ Standby MDM: A standby MDM can be called on to assume the position of a manager MDM when it is promoted to be a cluster member. ○ Standby Tiebreaker: A standby node that is prepared to take over as a tie breaker. <ul style="list-style-type: none"> ● Fault set: A logical group of SDSs within a protection domain that defines by the way it is grouped where the copies of data exist.

Edit service information

You can edit information about a service such as the service name, description, and who has access to the service information.

Steps

1. On the menu bar, click **Services**.
2. On the **Services** page, click the service, and in the right pane of the **Service Details** page, click **View Details**.
3. On the **Service Details** page, in the right pane click **Edit**.
4. In the **Edit Service Information** dialog box, modify the **Service Name** and **Service Description** that identifies the service.
 - To update the firmware and software running on the nodes when you deploy a service that uses this template, select **Firmware and Software Compliance**.

If you select a minimal compliance version for a service, PowerFlex Manager puts the service in lifecycle mode and restricts the actions that can be performed. In lifecycle mode, the service supports monitoring, service mode, and compliance upgrade operations only. All other service operations are blocked.

 **NOTE:** Changing the firmware repository might update the firmware level on nodes for this service. The global default firmware repository maintains firmware on shared devices.

5. Set the **OS Image** for the service.

You can optionally select a custom operating system image that will be used for future add node operations on the service. The custom image must be added to the **OS Image Repositories** tab on the **Settings > Compliance and OS Repositories** page.

PowerFlex Manager filters the list of images available for selection to include only those images that are suitable for the current service. For example, if you are editing a hyperconverged service, PowerFlex Manager only allows you to choose ESXi images.

6. Specify the permissions for this service under **Who should have access to the service deployed from this template?** by performing one of the following actions:
 - To restrict access to administrators, select the **Only PowerFlex Manager Administrators** option.
 - To grant access to administrators and specific standard and operator users, select the **PowerFlex Manager Administrators and Specific Standard and Operator Users** option, and perform the following tasks:
 - a. Click **Add User(s)** to add one more standard or operator users to the list.
 - b. Select which users will have access to this service.
 - c. To delete the standard and or operator user from the list, select the user and click **Remove User(s)**.
 - d. After adding the standard and or operator users, select or clear the check box next to the standard or operator users to grant or block access to use this template.
 - To grant access to administrators and all standard and operator users, select the **PowerFlex Manager Administrators and All Standard and Operator Users** option.
7. Click **Save**.

Service actions

You can perform various actions to a service that is deployed.

The following table describes the specific tasks that you can perform from the **Service Actions** section on the **Service Details** page.

Click...	To...
Retry	Redeploy a failed or canceled service. This button is visible only when a service is in a critical state.
Update Service Details	Update service definition.
Remove Service	Remove a service that is no longer required. PowerFlex Manager supports two types of removal: <ul style="list-style-type: none"> • Delete Service: Deletes the deployment information and makes any required configuration changes for components that are associated with the service. The configuration changes include making destructive changes such as unconfiguring the nodes in the service and powering them down. • Remove Service: Removes deployment information but does not make any configuration changes to the deployed components.
Cancel	Cancel the service.
View All Settings	View the settings configured on the resources in a service for deployment. For a PowerFlex service, the VMware Cluster section of the Settings page shows the PowerFlex Settings section, and the Target PowerFlex Gateway should specify the gateway that is included with the service. For a PowerFlex service, the Use Local Storage For Dell EMC PowerFlex property is set to true within the OS Settings section for any node included. This property is set to true for the nodes in an existing PowerFlex service deployment, and for the nodes in a new PowerFlex service deployment.
Generate troubleshooting bundle	Generate a compressed file of logs to use for troubleshooting. The following logs are included: <ul style="list-style-type: none"> • PowerFlex Manager application • Alert connector • PowerFlex Gateway • iDRAC life cycle • Dell EMC Powerswitch switch • Cisco Nexus switch

Click...	To...
	<ul style="list-style-type: none"> VMware ESXi CloudLink Center
Enter Service Mode	Put nodes in service mode so that maintenance operations can be performed.
Exit Service Mode	Take nodes out of service mode.
View Compliance Report	View the software or firmware compliance report.
Reconfigure MDM Roles	Change the MDM role for a node in a PowerFlex cluster. For example, if you add a node to the cluster, you might want to switch the MDM role from another node to the new node.

Service mode options

PowerFlex Manager supports short-term and long-term maintenance options for putting a node in service mode.

PowerFlex Manager detects when a node is in VMware ESXi or PowerFlex maintenance mode. It automatically places the node in service mode and also ensures that the service itself goes into service mode.

Mode	Description
Instant maintenance mode	Enables you to perform updates quickly. PowerFlex Manager does not migrate the data.
Protected maintenance mode	Enables you to perform updates that require longer than 30 minutes, in a safe and protected manner. When you use protected maintenance mode, PowerFlex makes a temporary copy of the data so that the cluster is fully protected from data loss. Protected maintenance mode applies only to PowerFlex hyperconverged and storage-only nodes.
Evacuate node from PowerFlex	Available if using a PowerFlex version prior to 3.5. Enables you to perform long-term maintenance that lasts more than 30 minutes. PowerFlex Manager migrates the data to other nodes in the cluster. Evacuate mode applies only to PowerFlex hyperconverged and storage-only nodes.

Enter and exit service mode

Place a node into service mode to perform maintenance operations on the node.

About this task

When you put a node into service mode, you specify whether you are performing short-term or long-term maintenance work. The option that you use for long-term maintenance depends on the PowerFlex version you are using.

PowerFlex Manager does not allow a node to enter service mode in the following scenarios:

- DAS Cache is installed.
- VMware NSX-T or NSX-V is configured.
- PowerFlex Gateway used in the service is being updated on the **Resources** page.
- Services that do not have switches discovered.

Prerequisites

Before you evacuate a node for long-term maintenance work, ensure:

- There are at least four nodes in the cluster.
- There is enough storage space on the remaining nodes so that the data from the evacuated node can be migrated to one of the remaining nodes.

If using protected maintenance mode, the sum of the spare and free capacity must be greater than the size of the node being put into protected maintenance mode.

Steps

1. On the menu bar, click **Services**.
2. On the **Services** page, select a service and click **View Details** in the right pane.
3. Click **Enter Service Mode** under **Service Actions**.
4. Select one or more nodes on the **Node Lists** page and click **Next**.

You can put multiple nodes in service mode simultaneously only if all the nodes are in the same fault set.

5. Specify the type of maintenance you want to perform by selecting one of the following options:
 - **Instant Maintenance Mode:** Use to perform short-term maintenance that lasts less than 30 minutes. PowerFlex Manager does not migrate the data.
 - **Protected Maintenance Mode:** Use to perform maintenance that requires longer than 30 minutes in a safe and protected manner. When you use protected maintenance mode, PowerFlex makes a temporary copy of the data so that the cluster is fully protected from data loss. Protected maintenance mode applies only to hyperconverged and storage-only services.
 - **Evacuate Node from PowerFlex** (PowerFlex version prior to 3.5): Use to perform long-term maintenance that lasts more than 30 minutes. PowerFlex Manager migrates the data to other nodes in the cluster. It takes longer to evacuate a node, but it is safer because there is no risk of a reboot causing data to be unavailable. Evacuation mode applies only to hyperconverged and storage-only services.
6. Click **Enter Service Mode**.

PowerFlex Manager displays a yellow warning banner at the top of the service page. The **Service Mode** icon displays for the **Deployment State** and **Overall Service Health**, and for the **Resource Health** for the selected nodes.
7. To leave service mode, click **Service Actions > Exit Service Mode**.

Adding volumes

Add a volume to a PowerFlex hyperconverged node

Initially, there are no volumes for a new service. Add at least one volume to a service for it to be in a complete state. If the service is a VMware ESXi cluster that has VMware High Availability (HA) enabled, you must add at least two volumes.

About this task

For a PowerFlex hyperconverged node, you must enter the datastore name for each new volume that must be added. PowerFlex Manager allows you to enable compression when adding volumes to a PowerFlex hyperconverged node.

Steps

1. On the **Service Details** page, click **Add Resources > Add Volumes**.
2. In the **Add Volume** wizard, click **Add volume**.
3. In the **Volume Name** field, select **Create New Volume** to create a volume now, or select an existing volume.
4. In the **New Volume Name** field, enter the volume name, if you are creating a volume.
5. In the **Datastore Name** field, select **Create New Datastore** to create a datastore or select an existing datastore.
6. In the **New Datastore Name** field, enter the datastore name, if you are creating a datastore.
7. From the **Storage Pool** menu, choose the storage pool where the volume resides.
8. Select the **Enable Compression** check box to use PowerFlex NVDIMM compression.
9. In the **Volume Size (GB)** field, select the size in GB. The minimum size is 8 GB. You must specify a value divisible by eight.
10. In the **Volume Type** field, select thick.
11. Optionally, click **Add volume** to add another volume. Provide the required information for the volume.
12. Click **Save**.

The service moves to the **In Progress** state and the new volume icons appear on the **Service Details** page. After the deployment completes, the new volumes appear on the **Service Details** page with a check mark in the **Storage** list.

Add a volume to a PowerFlex storage-only node

Initially, there are no volumes for a new service. Add at least one volume to a service for it to be in a complete state. If the service is an VMware ESXi cluster that has VMware High Availability (HA) enabled, you need to add at least two volumes.

About this task

When you add volumes to a PowerFlex storage-only node, the volumes are added to PowerFlex but are not mapped.

PowerFlex Manager allows you to enable compression when adding volumes to a PowerFlex storage-only node.

Steps

1. On the **Service Details** page, click **Add Resources > Add Volumes**.
2. In the **Add Volume** wizard, click **Add volume**.
3. In the **Volume Name** field, select **Create New Volume** to create a volume now.
4. In the **New Volume Name** field, type the volume name.
5. From the **Storage Pool** menu, choose the storage pool where the volume resides.
6. Select the **Enable Compression** check box to use PowerFlex NVDIMM compression.
7. In the **Volume Size (GB)** field, select the size in GB. The minimum size is 8 GB. You must specify a value divisible by eight.
8. In the **Volume Type** field, select thick or thin. If you enable compression, thin is the only available option.
9. Optionally, click **Add volume** to add another volume. Provide the required information for the volume.
10. Click **Save**.

The service moves to the **In Progress** state and the new volume icons appear on the **Service Details** page. After the deployment completes, the new volumes appear on the **Service Details** page with a check mark in the **Storage** list.

Add a volume to a PowerFlex compute-only node

Initially, there are no volumes for a new service. Add at least one volume to a service for it to be in a complete state. If the service is an VMware ESXi cluster that has VMware High Availability (HA) enabled, you need to add at least two volumes.

About this task

When you add volumes to a PowerFlex compute-only node, PowerFlex maps the volumes and creates the datastores.

Steps

1. On the **Service Details** page, click **Add Resources > Add Volumes**.
2. In the **Add Volume** wizard, click **Add volume**.
3. In the **Volume Name** field, select an existing volume. You can only select an existing volume that has not yet been mapped.
4. In the **Datastore Name** field, select **Create New Datastore** to create a datastore, or select an existing datastore. The **Datastore Name** field is only available for a PowerFlex hyperconverged node or PowerFlex compute-only node, as it applies only to services with VMware ESXi. If the volume was originally created on a PowerFlex storage-only node, you must select **Create New Datastore** to create a datastore. Alternatively, if the volume was originally created in a PowerFlex hyperconverged node, you must select the datastore that was already mapped to the selected volume in the other service.
5. In the **New Datastore Name** field, type the datastore name, if you are creating a datastore.
6. Optionally, click **Add volume** to add another volume. Provide the required information for the volume.
7. Click **Save**.

The service moves to the **In Progress** state and the new volume icons appear on the **Service Details** page. After the deployment completes, the new volumes appear on the **Service Details** page with a check mark in the **Storage** list.

Resize a volume

After adding volumes to a service, you can resize the volumes.

About this task

For a storage-only service, you can increase the volume size. For a VMware ESXi compute-only service, you can increase the size of the datastore that is associated with the volume. For a hyperconverged service, you can increase the size of both the volume and the datastore.

If you resize a volume in a storage-only service, you must update the datastore size in the corresponding VMware ESXi compute-only service. The datastore size cannot exceed the size of the volume.

Steps

1. On the **Services** page, click the volume that you want to resize and choose **Volume Actions > Resize**.
2. When PowerFlex Manager displays the **Resize Volume** wizard, update the sizing information.

If you are resizing a volume for a hyperconverged service, perform these steps:

- a. In the **New Volume Size (GB)** field, specify a value that is greater than the current volume size.
- b. Optionally, select **Resize Datastore** to increase the size of the datastore.

If you are resizing a volume for a storage-only service, enter a value in the **New Volume Size (GB)** field. Specify a value that is greater than the current volume size. Values need to be in multiple of eight or an error will occur.

If you are resizing a volume for a compute-only service, review the **Volume Size (GB)** field to see if the volume size is greater than **Current Datastore Size (GB)**. If it is, PowerFlex Manager expands the datastore size.

3. Click **Save**.

Add a network to a service

You can add an available network to a service or choose to define a new network for a configuration that was initially deployed outside of PowerFlex Manager. You cannot remove an added network using PowerFlex Manager.

About this task

Before you can add a network to a service, define the network.

You can add a static route to allow nodes to communicate across different networks. The static route can also be used to support replication.

Prerequisites

Ensure that a new VLAN is created on any switches that need access to that VLAN and is added to any management cluster server-facing ports. The VLAN is then added to any northbound trunks to other switches that it must communicate with.

Steps

1. On the menu bar, click **Services**.
2. Select a service for which you want to add a network and in the right pane, click **View Details**.
3. Under **Resource Action**, from the **Add Resources** list, click **Add Network**.
The **Add Network** window is displayed. All used resources and networks are displayed under **Resource Name** and **Networks**.
4. Click **Add Additional Network** to add an additional network:
 - a. From the **Available Networks** list, select the network, and click **Add**.
The selected network is displayed under **Network Name**. You can define a new network by clicking **Define a New Network**.
 - b. Select **Port Group** from the **Select Port Group** list.
 - c. Click **Save**.
5. Click **Add Additional Static Route** to add an additional static route:

- a. Click **Add New Static Route**.
- b. Select a **Source Network**.
The source network must be a PowerFlex data network or a replication network.
- c. Select a **Destination Network**.
The destination network must be a PowerFlex data network or a replication network.
- d. Type the IP address for the **Gateway**.
- e. Click **Save**.

Reconfiguring MDM roles

PowerFlex Manager enables you to change the MDM role for a node in a PowerFlex cluster. For example, if you add a node to the cluster, you might want to switch the MDM role from another node to the new node.

About this task

You can launch the wizard for reconfiguring MDM roles from the **Services** page or from the **Resources** page. The nodes that are listed and the operations available are the same regardless of where you launch the wizard.

Each fault set can have a maximum of one MDM role (management or tiebreaker). PowerFlex Manager blocks your role assignments if the reconfigured roles do not conform to PowerFlex best practices.

Steps

1. To access the wizard from the **Services** page:
 - a. On the menu bar, click **Services**.
 - b. Select a service that has the PowerFlex Gateway for which you want to reconfigure MDM roles.
 - c. In the right pane, click **View Details**.
The **Service Details** page is displayed.
 - d. On the **Service Details** page, click **Reconfigure MDM Roles**.
You can also click **Reconfigure MDM Roles** on the **Node Actions** menu on the **Services Details** page.
The **MDM Reconfiguration** page is displayed.
2. To access the wizard from the **Resources** page:
 - a. On the menu bar, click **Resources**.
 - b. Select the PowerFlex Gateway for which you want to reconfigure MDM roles.
 - c. In the right pane, click **View Details**.
The **Details** page is displayed for the selected PowerFlex Gateway.
 - d. On the **Details** page, click **Reconfigure MDM Roles**.
The **MDM Reconfiguration** page is displayed.
3. Review the current MDM configuration for the cluster.
4. Find each MDM role that you want to reassign and choose the new Host Name or IP address for the role in the **Select New Node for MDM Role** drop-down.
You can reassign multiple roles at one time.
5. Click **Next**.
The **Summary** page is displayed.
6. Type **CHANGE MDM ROLES** to confirm your changes.
7. Click **Finish**.

Remove a service

You can remove a service that is no longer required.

About this task

PowerFlex Manager supports the removal of an entire service or the removal of the deployment information from a service.

Option	Description
Delete service	Deletes deployment information and makes configuration changes to the components associated with the service to unconfigure those components. It returns the components to the available inventory. Before a rebuild, ensure that residual components, for example a data center in VMware vCentre, are removed by doing a manual cleanup.
Remove service	Removes deployment information without making any configuration changes to the deployed components. It returns the components to the available inventory. Use the remove service option if you want to change all the IP addresses or if new networks were added outside of PowerFlex Manager. The service is not removed permanently and can be added back as an existing service to PowerFlex Manager.

If DAS Cache is installed on a node, or the node has a VMware NSX-T or NSX-V configuration, you can remove the deployment information for a service. However, the service is not deleted. If the PowerFlex Gateway used in a service is being updated on the **Resources** page, PowerFlex Manager does not allow you to delete the service.

Standard users can delete only the services that they have deployed.

Steps

1. On the menu bar, click **Services**.
2. On the **Services** page, click the service, and in the right pane, click **View Details**.
3. On the **Service Details** page, in the right pane, under **Service Actions**, click **Remove Service**.
4. In the **Remove Service** dialog box, select the **Service removal type**.
5. If you choose **Delete Service**, perform the following steps:
 - a. Select **Delete Clusters(s) and Remove from vCenter** to delete and remove the clusters from vCenter.
 - b. Select **Remove Protection Domain and Storage Pools from PowerFlex** to remove the protection domain and storage pools that are created during the service deployment.

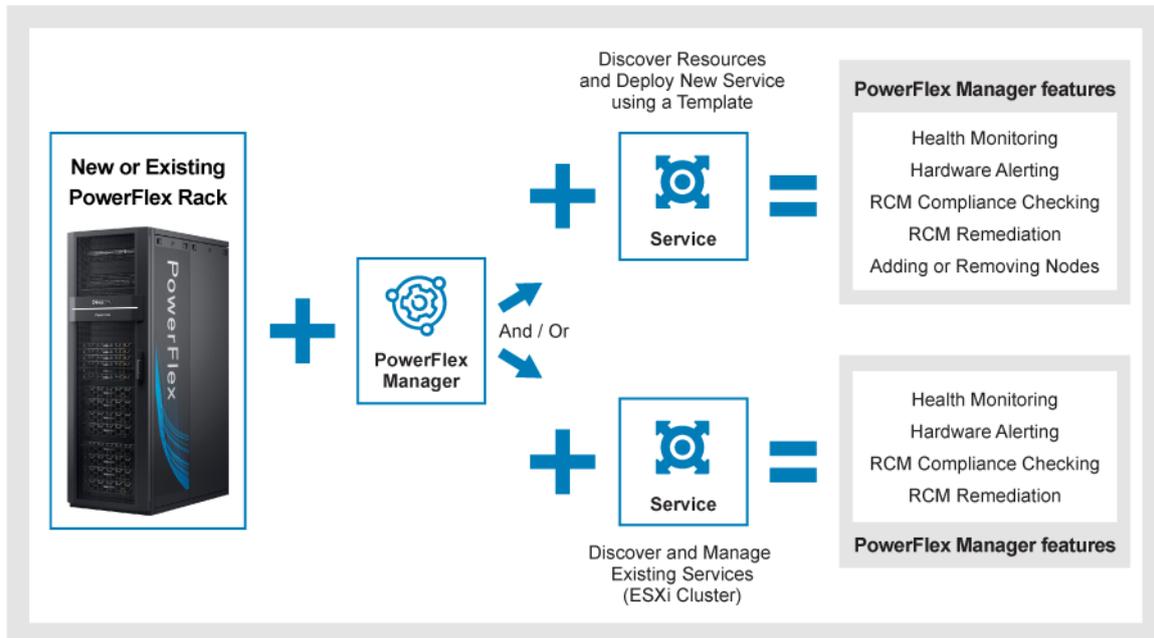
If you select this option, you must select the target PowerFlex Gateway. The PowerFlex Gateway is not removed. PowerFlex Manager removes only the protection domain and storage pools that are part of the service. If multiple services are sharing a protection domain, you might not want to delete the protection domain.

For a compression enabled service, PowerFlex Manager deletes the acceleration pool and the DAX devices when you delete the service.
 - c. Select **Delete Machine Group and remove from CloudLink Center** to clean up the related components in CloudLink Center.

CloudLink Center cleanup includes the deletion of the machine group, keystore, and approved network related to the service being deleted. These components are removed from CloudLink Center only if all the machines that are associated to the machine group are deleted first. If all the machines that are related to this machine group are not removed, the cleanup does not succeed as there are machines associated with the machine group.
 - d. If you are certain that you want to proceed, type **DELETE SERVICE**.
 - e. Click **Delete**.
6. If you choose **Remove Service**, perform the following steps:
 - a. To keep the nodes in the inventory, select **Leave nodes in PowerFlex Manager inventory and set state to** and select the state of managed, unmanaged, or reserved.
 - b. To remove the nodes, select **Remove nodes from the PowerFlex Manager inventory**.
 - c. Click **Remove**.

Services deployed outside of PowerFlex Manager

PowerFlex Manager can discover PowerFlex rack environments that are deployed outside of PowerFlex Manager. It discovers existing VMware vSphere clusters as existing services. PowerFlex Manager can then perform alerting, health monitoring, and RCM compliance monitoring and remediation on these services.



During discovery, only resources with a status of managed or reserved are in the existing service. Unmanaged resources are not included. If you find a component missing, you can change its state to managed and update the service details to include it.

The existing service must fall within the hardware compatibility matrix and be supported at current minimum RCM levels.

Keep the following considerations in mind when discovering existing services:

Component	Details
Network switch	<p>You can discover an existing service without switches. PowerFlex Manager places the service in lifecycle mode, which limits the actions you can perform within the service. Lifecycle mode allows the service to perform only monitoring, service mode, and non-disruptive compliance upgrade operations.</p> <p>The switches must be manually configured correctly to prevent PowerFlex Manager actions requiring switch connectivity from failing.</p>
CloudLink Center	<p>When you add an existing service for a CloudLink configuration, you select a target CloudLink Center for the service. If PowerFlex Manager loses communication with the CloudLink Center and it is part of a cluster, PowerFlex Manager moves to another CloudLink Center when you update service details.</p>
Nodes	<p>You can select one credential for all nodes or choose a separate credential for each node. Go to the Credentials Management page under Settings to manage credentials. You can select one credential for all nodes (or SVMs) or choose a separate credential for each node. Go to the Credentials Management page under Settings to manage credentials. PowerFlex Manager inventories all nodes with valid credentials inventories all nodes with valid credentials.</p>
Network	<p>PowerFlex Manager attempts to select the correct network that is based on the information provided. If it finds more than one network, you must select the network from the available list.</p>
Port group	<p>If you do not want PowerFlex Manager to manage access for a port group, leave it cleared. A network must be selected for a port group to be in the deployment.</p>
PowerFlex management controller	<p>You cannot discover a PowerFlex management controller as an existing service.</p>

You can distinguish existing services from new services on the **Service Details** page. For an existing service, the **Reference Template** field contains **Generated Existing Service Template**.

Add a service deployed outside of PowerFlex Manager

You can add an existing service to discover and import hardware resources that were not originally deployed with PowerFlex Manager. Most of the actions that can be performed within PowerFlex Manager for a new service can also be performed on a service that was deployed outside of PowerFlex Manager.

About this task

You cannot add a PowerFlex management controller.

Prerequisites

- If you are using node pools, you must have access to the node pool.
- The VMware vCenter, PowerFlex Gateway, CloudLink Center, and hosts must be discovered in the resource list.
- The PowerFlex Gateway must be in the service.

Steps

1. On the menu bar, click **Services > + Add Existing Service**.
2. On the **Add Existing Service** page, enter a service name in the **Name** field.
3. Enter a description in the **Description** field.
4. Select the **Type** for the service.
5. To specify the version to use for compliance, select the version from the **Firmware and Software Compliance** list or choose **Use PowerFlex Manager appliance default catalog**.
 **NOTE:** Changing the compliance version might update the firmware level on nodes for this service. The global default firmware repository maintains the firmware on shared devices.
6. Specify the service permissions under **Who should have access to the service deployed from this template?**. Click **Next**.
7. Choose the network automation type. If you select partial network automation, PowerFlex Manager skips the switch configuration step that is normally performed for a service with full network automation. Partial network automation enables you to work with unsupported switches. It requires more manual configuration before a deployment can successfully proceed.
 In the **Number of Instances** box, provide the number of component instances to include in the template.
8. On the **Cluster Information** page, enter a name for the cluster component. Select values for the cluster settings:

Type	Cluster settings	Description
Hyperconverged or compute only	Target Virtual Machine Manager	Select the vCenter name where the cluster is available.
	Data Center Name	Select the data center name where the cluster is available.  NOTE: Ensure that the selected vCenter has unique names for clusters in case there are multiple clusters in the vCenter.
	Cluster Name	Select the name of the cluster to discover.
	OS Image	Select the image or choose Use Compliance File ESXi image if you want to use the image provided with the target compliance version. PowerFlex Manager filters the operating system image choices to show only ESXi images for a hyperconverged or compute-only service.
Storage only	Target PowerFlex Gateway	Select the gateway where the cluster is available.
	Protection Domain	Select the name of the protection domain in PowerFlex.

Type	Cluster settings	Description
	OS Image	Select the image or choose Use Compliance File Linux image to use the image provided with the target compliance version. PowerFlex Manager filters the operating system image choices to show only Linux images for a storage-only service.

9. Click **Next**.
10. On the **OS Credentials** page, select the OS credential that you want to use for each node and SVM. You can select one credential for all nodes (or SMVs) or choose credentials for each item separately. Click **Next**.
11. Review the inventory on the **Inventory Summary** page and click **Next**.
12. On the **Network Mapping** page, review the networks that are mapped to port groups, make any required edits, and click **Next**.
13. To import a large number of general purpose VLANs from VMware vCenter, perform these steps:
 - a. On the **Network Mapping** page, click **Import Networks**. PowerFlex Manager lists as **Available Networks** the port groups that are defined in the vCenter. You can see the port groups and VLAN IDs.
 - b. Optionally, search for a VLAN name or VLAN ID to limit the search results.
 - c. Click each network that you want to add under **Available Networks**. To add all available networks, click the check box to the left of the **Name** column.
 - d. Click the double arrow (>>) to move the networks to **Selected Networks**.
 - e. Click **Save**.
14. Review the **Summary** page and click **Finish** when you are ready to add the service.

Update resource inventory

Use this procedure to run inventory to incorporate changes that are made to resource data outside of PowerFlex Manager.

About this task

NOTE: Run inventory to incorporate external changes that are made to resource data outside of PowerFlex Manager. After running the inventory to incorporate these changes, you can update the details on any service that must include the new resource data.

Administrators and standard users can run the inventory on any resources. Standard users can run the inventory only on resources that are part of a node pool for which they have permission.

When you run the inventory, PowerFlex Manager checks to see whether DAS Cache is installed on any of the SVMs. If DAS Cache is installed, PowerFlex Manager disables some service features. PowerFlex Manager does not enable you to expand or remove resources within a service. You cannot enter service mode, perform a drive replacement, perform firmware updates, delete a service, or retry a service deployment.

Steps

1. On the menu bar, click **Resources**.
2. On the **Resources** page, click the **All Resources** tab.
3. From the list of resources, select the check box next to the resources that you want to inventory.
4. From the **Details** pane, click **Run Inventory**.

Next steps

See the PowerFlex Manager logs to view the start time and end time of the resource inventory operation. Go to **Settings > Logs**.

Discovering resources

You can discover new resources or existing resources that are already configured within your environment. After discovery, you can deploy services on these resources from a template. Only administrator-level users can discover resources.

By default, the operational state for all discovered nodes is **Unmanaged**. If you want to perform firmware updates or deployments on a discovered node, select the node and change the operational state to **Managed**.

For some resources such as nodes, the default credentials are prepopulated in PowerFlex Manager. If the credentials are changed from the defaults, add the credential to PowerFlex Manager with the new login information.

Element manager discovery

The IPI cabinet is discovered as an element manager. Once PowerFlex Manager discovers the IPI cabinet, you can view cabinet details or open the IPI appliance management application to view more details.

CloudLink Center is discovered as an element manager. If you create your own username in CloudLink Center, you need the SecAdmin access role and the local user type for the new user to complete discovery. The client user type does not work for discovery in PowerFlex Manager. If one member of a cluster of CloudLink Centers is shut down, the inventory still succeeds because PowerFlex Manager inventories the other CloudLink Center cluster members.

Node discovery

PowerFlex Manager supports PowerFlex node discovery and allows you to onboard nodes by configuring the initial management IP address and iDRAC credentials. To perform initial discovery and configuration, verify that the management IP address is set on the node and that PowerFlex Manager can access the IP address through the network. While configuring IP addresses on the node, verify that PowerFlex Manager can access any final IP address in a range used for hardware management, to complete discovery of these nodes.

PowerFlex Manager allow you to use name-based searches to discover a range of nodes that were assigned IP addresses by DHCP to iDRAC. You can search for a range of DNS hostnames or a single hostname within the Discovery wizard. After you perform a name-based discovery, PowerFlex Manager operations to iDRAC continue to use name-based IP address resolution.

Switch discovery

If you attempt to discover a Cisco switch with terminal color configured, the discovery fails. To discover the switch successfully, disable the terminal color option by running `configure no terminal color persist`.

Virtual machine manager discovery

PowerFlex Manager users with the administrator role can discover a VMware vCenter as a VM manager in PowerFlex Manager. A vCenter read-only user can discover a vCenter in PowerFlex Manager only after the following requirements are met:

- The vCenter user who is specified in the vCenter credential is granted the `VirtualMachine.Provisioning.ReadCustSpecs` and `StorageProfile.View` privileges.
- The permission containing these privileges is granted to that user on the root vCenter object and the `Propagate to children` property is set to `True`.

PowerFlex Manager allows you to deploy new hyperconverged and compute-only services and add existing services to a vCenter that has an enhanced link mode (ELM) configuration. ELM connects multiple vCenter servers, allowing you to search across servers and perform other vCenter management functions. ELM does not provide clustering or redundancy.

If you are working with a vCenter that has an enhanced link mode configuration, you must discover only the vCenter to which you want to deploy or add an existing service. You do not need to discover the other vCenters that are connected through the enhanced link mode configuration.

Discover a resource

A resource must be discovered in PowerFlex Manager in order for PowerFlex Manager to manage it. A resource in PowerFlex Manager is categorized into one of the following groups: element manager, node, switch, VM manager, or PowerFlex Gateway.

Prerequisites

Before you start discovering a resource, complete the following:

- Gather the IP addresses and credentials that are associated with the resources.
- Ensure that both the resources and the PowerFlex Manager virtual appliance are connected to the network.

Steps

1. Access the **Discovery Wizard** by performing either of the following actions:
 - On the **Getting Started** page, click **Discovery Resources**.
 - On the menu bar, click **Resources**. On the **Resources** page, click **Discover** on the **All Resources** tab.
2. On the **Welcome** page of the **Discovery Wizard**, read the instructions, and click **Next**.
3. On the **Identify Resources** page, click **Add Resource Type**, and perform the following steps:
 - a. From the **Resource Type** list, select a resource that you want to discover.
 - **Element Manager**, for example, CloudLink Center.
 - **Node**
 - **Switch**
 - **VM Manager**
 - **PowerFlex Gateway**
 - b. Enter the management IP address (or hostname) of the resources that you want to discover in the **IP/Hostname Range** field.

To discover one or more nodes by IP address, select **IP Address** and provide a starting and ending IP address.

To discover one or more nodes by hostname, select **Hostname** and identify the nodes to discover in one of the following ways:

- Enter the fully qualified domain name (FQDN) with a domain suffix.
- Enter the FQDN without a domain suffix.
- Enter a hostname search string that includes one of the following variables:

Variable	Description
\${num}	Produces an automatically generated unique number.
\${num_2d}	Produces an automatically generated unique number that has two digits.
\${num_3d}	Produces an automatically generated unique number that has three digits.

If you use a variable, you must provide a start number and end number for the hostname search.

- c. Select one of the following options from the **Resource State** list:

Option	Description
Managed	<p>Select this option to monitor the firmware version compliance, upgrade firmware, and deploy services on the discovered resources. A managed state is the default option for the switch, vCenter, element manager, and PowerFlex Gateway resource types.</p> <p>Resource state must be set to Managed for PowerFlex Manager to send alerts to Secure Remote Services.</p>

Option	Description
Unmanaged	Select this option to monitor the health status of a device and the firmware version compliance only. The discovered resources are not available for a firmware upgrade or deploying services by PowerFlex Manager. This is the default option for the node resource type. If you did not upload a license in the Initial Setup wizard, PowerFlex Manager is configured for monitoring and alerting only. In this case, Unmanaged is the only option available.
Reserved	Select this option to monitor firmware version compliance and upgrade firmware. The discovered resources are not available for deploying services by PowerFlex Manager.

- d. To discover resources into a selected node pool instead of the global pool (default), select an existing or create a node pool from the **Discover into Node Pool** list. To create a node pool, click the **+** sign to the right of the **Discover into Node Pool** box.
- e. Select an existing or create a credential from the **Credentials** list to discover resource types. To create a credential, click the **+** sign to the right of the **Credentials** box. PowerFlex Manager maps the credential type to the type of resource you are discovering. The credential types are as follows:
 - Element manager
 - Node
 - Switch
 - vCenter
 - PowerFlex Gateway

The default node credential type is **Dell EMC PowerEdge iDRAC Default**.

- f. If you want PowerFlex Manager to automatically reconfigure the iDRAC nodes it finds, select the **Reconfigure discovered nodes with new management IP and credentials** check box. This option is not selected by default, because it is faster to discover the nodes if you bypass the reconfiguration.
- g. To have PowerFlex Manager automatically configure iDRAC nodes to send alerts to PowerFlex Manager, select the **Auto configure nodes to send alerts to PowerFlex Manager** check box.

4. Click Next.

You might have to wait while PowerFlex Manager locates and displays all the resources that are connected to the managed networks.

To discover multiple resources with different IP address ranges, repeat step 2 and 3.

5. On the **Discovered Resources** page, select the resources from which you want to collect inventory data and click **Finish**. The discovered resources are listed on the **Resources** page.

Managing credentials

PowerFlex Manager requires a root-level username and password to access and manage PowerFlex nodes, switches, VMware vCenter, element managers, PowerFlex Gateway, and operating system resources.

Update a credential in PowerFlex Manager

If the password for a resource or component is manually changed outside of PowerFlex Manager, you must update the credential in PowerFlex Manager with the new password.

About this task

If you change the password for a resource, you must update all similar resources (for example, element manager, node, switches, VM Manager, PowerFlex Gateway) to have the same password.

Steps

1. Log in to PowerFlex Manager.
2. On the menu bar, click **Settings > Credentials Management**.
3. On the **Credential Management** page, select the resource whose password you want to edit, and click **Edit**.
4. In the **Edit Credentials** dialog box, modify the password.
5. Click **Save**.
6. On the menu bar, click **Services** and select the deployed service that contains the resource whose password is updated.
7. Under **Service Actions**, click **Update Service Details**.

Update a credential for a PowerFlex node in a service

When you view details for a service, PowerFlex Manager displays all PowerFlex nodes and credentials, regardless of whether they are in the service. If a PowerFlex node changes, update the username and password through the service actions.

Steps

1. On the menu bar, click **Services**.
2. On the **Services** page, click the service, and under **Service Actions**, click **Update Service Details**. The **Update Resources** page opens with the list of nodes to which you can assign credentials.
3. On the **OS Credentials** page, click **+**. The **Create Credentials** page opens.
4. Enter the required information in the **Credential Name**, **Password**, and **Confirm Password** fields.
5. Click **Save** and then click **Next**.
6. On the **Inventory Summary** page, review the information and click **Next**.
7. Review the **Summary** page, and click **Finish**.

Create a credential

You can create a credential on the **Credentials Management** page. Alternatively, you can create a credential while you are editing a template. If you select a credential that was created on the **Credentials Management** page, do not provide the username and password, since they are part of the credential definition.

About this task

The username **root** is the only valid username for root-level credentials on nodes (iDRAC). You can add iDRAC users with a username other than **root**.

For the OS user credential type, you can enter a user other than root. For the embedded operating system, this user account must have SSH enabled and have non-root access. For ESXi, the account must be configured with the administrator role on the local server permission setting, which enables SSH and other tools such as esxcli. You can add existing services with a non-root user.

For the OS admin credential type, the user name field is disabled because the user is assumed to be root. You must use the root user for new deployments.

Provide two separate usernames for the PowerFlex Gateway credential type: **Gateway Admin User Name** and **Gateway OS User Name**. The gateway admin user is the REST API administrator, and the username must be **admin**. The gateway operating system user is the SSH login user, and the username must be **root**.

Only a node or switch requires SNMP configuration.

Steps

1. On the menu bar, click **Settings > Credentials Management**.
2. On the **Credentials Management** page, click **Create**.
3. In the **Create Credentials** dialog box, from the **Credential Type** drop-down list, select one of the following resource types for which you want to create the credentials:
 - **Node**
 - **Switch**
 - **vCenter**
 - **Element Manager**
 - **PowerFlex Gateway**
 - **OS Admin**
 - **OS User**

The OS admin and OS user credential types apply to deployed items, not to the PowerFlex Manager virtual appliance itself.

4. In the **Credential Name** field, enter a name to identify the credential.
5. In the **User Name** field, enter a username for the credential.
6. In the **Password** and the **Confirm Password** boxes, enter the password for the credential.
7. To configure trap receiving for SNMPv2:
 - a. Under **SNMP Configuration**, select **V2** as the SNMP type.
 - b. Click **Add** next to the **SNMP v2 Community String** box.
The **SNMP v2 Community String** page opens.
 - c. Enter the community string by which PowerFlex Manager receives traps from devices and by which it forwards traps to destinations.
 - d. Click **Save**.

 **NOTE:** You can add more than one community string. For example, add more than one if the community string by which PowerFlex Manager receives traps differs from the community string by which it forwards traps to a remote destination.

8. To configure trap receiving for SNMPv3:
 - a. Under **SNMP Configuration**, select **V3** as the SNMP type.
 - b. Click **Add** beside the **SNMP V3 User** box.
The **SNMP V3 User** page opens.
 - c. Enter the **Username**, which identifies the ID where traps are forwarded on the network management system.

 **NOTE:** The username must be at most 16 characters.

d. Select a **Security Level** and provide a password if required.

You can add more than one user.

 **NOTE:** You can only select SNMPv3 for nodes. Select the **Auto configure nodes to send alerts to PowerFlex Manager** check box to discover iDRACs.

9. To save the credential, click **Save**.

Related information

[Configure SNMP](#) on page 49

Edit a credential

You can edit the name, username, and password associated with a credential.

Steps

1. On the menu bar, click **Settings > Credentials Management**.
2. On the **Credential Management** page, click a credential that you want to edit, and click **Edit**.
3. Modify the credential information in the **Edit Credentials** dialog box.
4. Click **Save**.

Delete a credential

To remove a credential that is part of a service, you must first remove the service.

Steps

1. On the menu bar, click **Settings > Credentials Management**.
2. On the **Credential Management** page, select the credential that you want to delete and click **Delete**.
3. Click **OK** when the confirmation message appears.

Updating passwords for system components

You can update the passwords for some system components from PowerFlex Manager.

Steps

1. On the menu bar, click **Resources**.
2. On the **All Resources** tab, select one or more resources of the same type for which you want to change passwords. For example, you could select one or more iDRAC nodes or you could select one or more PowerFlex Gateway components.
3. Click **Update Password**.
PowerFlex Manager displays the **Update Password** wizard.
4. On the **Select Components** page, select one or more components for which you want to update a password and click **Next**.
The component choices vary depending on which resource type you initially selected on the **Resources** page.
5. On the **Select Credentials** page, create a credential or change to a different credential having the same username. Click **Next**.
6. Click **Finish**.

Managing compliance

PowerFlex Manager compares the current firmware and software levels and target compliance versions, which contain the baseline firmware and software versions. Any deviation from the baseline in the compliance status of the resources is displayed. You can use PowerFlex Manager to initiate updates to bring the resources to a compliant state.

Compliance status

PowerFlex Manager assigns one of the following firmware statuses to the resources:

Firmware status	Description
Compliant	The firmware running on the resource is compliant with the firmware version that is specified in the default compliance version.
Non-Compliant	The firmware running on the resource is less than or greater than the firmware version specified in the default compliance version. Indicates that firmware update is required.
Update Required	The firmware running on the resource is less than the minimum firmware version recommended in the default compliance version. Indicates that firmware update is required.

Update firmware and software

You can update the firmware and software for a resource that is associated with a PowerFlex Gateway, switch, or PowerFlex node as long as it is not part of a service. If the firmware or software is not compliant with the default compliance catalog, you can update the firmware or software to the minimum recommended level.

About this task

If the PowerFlex Gateway, switch, or PowerFlex node is part of a service, see the [Dell EMC PowerFlex Rack Upgrade Guide](#) for more information.

A firmware update on a VMware cluster is successful only if the cluster is properly configured for HA. For example, the host system can be set to maintenance mode and the virtual machines can be moved from one host to another in the cluster.

A firmware or software update on a node that is part of a cluster is successful only if the node is set to maintenance mode. PowerFlex Manager sets nodes in a cluster to maintenance mode before performing an update. To ensure that the node remains in maintenance mode, ensure that there are other nodes available in the cluster to host the virtual machines of the node being updated.

If you use the **Resources** pane to update a resource that is part of a cluster, which is not added to PowerFlex Manager, the PowerFlex node will not be put into maintenance mode. If the host is not manually placed in maintenance mode, an outage occurs.

Steps

1. On the menu bar, click **Resources**.
2. Select the check box for one or more resources and click **Update Resources**.
3. On the **Update Details** page, confirm that all the information is correct. For a PowerFlex Gateway, reconfigure one or more nodes before proceeding with the update.

4. Click **Next**.
5. On the **Summary** page, select one of the following options:
 - **Allow PowerFlex Manager to perform firmware and software updates now**—Select this option to update the firmware and software immediately. PowerFlex Manager applies the firmware updates and reboots the selected resources immediately. If the resource is in use, this update could be disruptive.
 - **Schedule firmware and software updates**—Select this option and indicate the date and time to update the firmware and software.

PowerFlex Manager applies the firmware updates at the selected date and time and reboots the selected resources. If the resource is in use, this update could be disruptive.
6. If you are updating a PowerFlex Gateway, type **UPDATE POWERFLEX** to confirm that you are ready to proceed with the update.
7. Click **Finish**.

View a compliance report and perform remediation

PowerFlex Manager enables you to view a compliance report and perform firmware and software updates as needed.

About this task

If DAS Cache is installed on a node, you cannot update software and firmware components for the service.

The PowerFlex Gateway must first be compliant before you update a node. If a node is already compliant, it is disabled and cannot be selected.

If you want to update only the firmware on the nodes, select a minimal compliance version with firmware only as the target version for the service.

PowerFlex Manager does not upgrade VMware vCenter itself. It checks the vCenter version to determine if it matches the VMware ESXi version. If the ESXi version is higher, PowerFlex Manager blocks the ESXi host upgrade and displays an error. In this case, upgrade vCenter first or use a different compliance version that is compatible with the installed vCenter version.

Steps

1. On the menu bar, click **Services**.
2. Select a service to view the compliance report.
3. In the right pane, click **View Details**.
4. On the **Service Details** page, under **Service Actions**, click **View Compliance Report**.
5. Click **Firmware Components** to view the firmware components.
6. Click **Software Components** to view the software components.
 - For compute-only and storage-only services, the **Components** drop-down list shows the software components available for update.
 - For hyperconverged services, PowerFlex Manager displays the following categories of software components available for update: **SVM Software** and **Node Software**.
7. Select one or more noncompliant nodes. If you are planning to perform an upgrade for a full PowerFlex rack, click the box in the upper-left corner of the **Firmware Components** tab to select all nodes in the service.
8. Click **Update Resources**.
9. Specify when to apply the resource updates.
 - To perform a non-disruptive update right away, choose **Allow PowerFlex Manager to perform firmware and software updates now**. Specify the type of update you want to perform by selecting either **Instant Maintenance Mode**, **Protected Maintenance Mode**, or **Evacuate Node from PowerFlex** (for PowerFlex versions prior to 3.5).
 - To perform a non-disruptive update at a later time, choose **Schedule firmware and software updates**. Specify the type of update you want to perform by selecting either **Instant Maintenance Mode**, **Protected Maintenance Mode**, or **Evacuate Node from PowerFlex** (for PowerFlex versions prior to 3.5).
 - To perform a disruptive update right away for a full PowerFlex rack, choose **Allow PowerFlex Manager to perform disruptive updates now**. If you are certain you want to proceed, type **REBOOT ALL NODES AT ONCE**.
10. Click **Apply**.

Next steps

If you encounter any errors while performing firmware or software updates, you can view the PowerFlex Manager logs for the service. On the **Service Details** page, in the right pane under **Service Actions**, click **Generate Troubleshooting Bundle**.

Export a compliance report for all resources

You can download a CSV or PDF report that lists compliance details for all resources. The CSV report includes all information and can be imported into a database for querying. The PDF contains a subset of information to make it easier to read.

About this task

Each resource is measured against a specified compliance file version. For each resource, the report lists the components, the current and expected software or firmware version, and whether the component is compliant.

Steps

1. On the menu bar, click **Resources**.
2. From the **Export Report** drop-down list, select either **Export Compliance PDF Report** or **Export Compliance CSV Report**.
The selected report downloads.

Updating PowerFlex Manager and components

Update the PowerFlex Manager virtual appliance before you update a component. After you update the virtual appliance, a notification banner might display at the top of the **Services** page. The banner alerts you that a service is in a critical state.

To get the service into a healthy state, you must update the components if PowerFlex Manager added new required fields to any of the components.

Use the PowerFlex Manager version applicable to the RCM. If you must upgrade the PowerFlex Manager virtual appliance to the supported latest version, PowerFlex Manager has an integrated mechanism for upgrading the PowerFlex Manager virtual appliance from an older version to the latest version.

See the [Dell EMC PowerFlex Rack Upgrade Guide](#) for information about updating the PowerFlex Manager virtual appliance and incrementally upgrading PowerFlex rack software from one RCM to the next.

Expanding a system

You can expand a system with additional PowerFlex nodes. You do this in PowerFlex Manager by adding nodes to a service.

The following are the high-level expansion tasks, if PowerFlex Manager was installed and configured in the factory:

1. Discover and configure the PowerFlex node for alerting.
2. Ensure the availability of network IP addresses.
3. Add PowerFlex nodes to a service.

If your system did not have PowerFlex Manager installed and configured in the factory, and you are expanding a service that was initially deployed outside of PowerFlex Manager, you must complete the following expansion steps:

1. Discover and configure the PowerFlex node for alerting.
2. Add a network to a service deployed outside of PowerFlex Manager.
3. Ensure the availability of network IP addresses.
4. Add PowerFlex nodes to a service.

Discover and configure a PowerFlex node for alerting

After the physical PowerFlex node is cabled and connected, PowerFlex Manager discovers and configures the PowerFlex node for alerting.

About this task

During PowerFlex node discovery, you can configure iDRAC nodes to automatically send alerts to PowerFlex Manager. If the PowerFlex nodes are not configured for the alert connector, Secure Remote Services does not receive critical or error alerts for those resources.

Prerequisites

Gather the IP addresses and credentials that are associated with the resources.

Steps

1. On the menu bar, click **Resources**.
 2. On the **Resources** page, click **Discover** on the **All Resources** tab.
 3. On the **Welcome** page of the **Discovery** wizard, read the instructions and click **Next**.
 4. On the **Identify Resources** page, click **Add Resource Type**. From the **Resource Type** list, select **Node**.
 5. Enter the management IP address of the node in the **IP Address Range** field.
 6. In the **Resource State** list, select **Managed**. The resource state must be managed for PowerFlex Manager to send alerts to Secure Remote Services.
 7. If expanding into a node pool, select the node pool from the **Discover into Node Pool** list.
 8. Select the appropriate iDRAC credential from the **Credentials** list.
 9. If you want PowerFlex Manager to automatically reconfigure the iDRAC nodes it finds, select the **Reconfigure discovered nodes with new management IP and credentials** check box. This option is not selected by default, because it is faster to discover the nodes if you bypass the reconfiguration.
-  **NOTE:** If you select this option, the iDRAC IP address and or password is changed.
10. Select the **Auto configure nodes to send alerts to PowerFlex Manager** check box to have PowerFlex Manager automatically configure iDRAC nodes to send alerts to PowerFlex Manager.
 11. Click **Next** to start discovery. On the **Discovered Resources** page, select the resources from which you want to collect inventory data and click **Finish**.

Related information

[Add a service deployed outside of PowerFlex Manager](#) on page 33

Ensure the availability of network IP addresses

When you expand your system, you must ensure that there are network IP addresses available for communication.

About this task

The IP address ranges cannot overlap and must be used exclusively for the static network defined. For example, you cannot create an IP address range of 10.10.10.1–10.10.10.100 and another range of 10.10.10.50–10.10.10.150.

Node type	IP address requirements
PowerFlex hyperconverged node	One VMware ESXi management IP address
	One vMotion IP address
	Two IP addresses per data network
	One PowerFlex management IP address
PowerFlex storage-only node	One IP address per data network
	One PowerFlex management IP address
PowerFlex compute-only node	One ESXi management IP address
	One vMotion IP address
	One IP address per data network
	One PowerFlex management IP address

Steps

1. Go to **Settings > Networks**.
2. Select the network on which you want to ensure that there are available IP addresses.
3. Scroll to the **Static IP Address Details** field, and select **Available IP Addresses** from the **View** drop-down list.

Add PowerFlex nodes to a service

After the additional PowerFlex node is discovered, you must add it to a service for it to be managed within PowerFlex Manager.

Steps

1. On the menu bar, click **Services**.
2. Select the service to which you are adding a PowerFlex node to view its details.
3. On the **Service Details** page, click **Add Resources**. Choose **Add Nodes**.
4. In the **Duplicate Node** wizard:
 - a. From the **Resource to Duplicate** list, select the node.
 - b. In the **Number of Instances** box, set the number of instances to **1** and click **Continue**.
 - c. Under **PowerFlex Settings**, specify the **PowerFlex Storage Pool Spare Capacity** setting.
 - d. Under **OS Settings**, set the **Host Name Selection**. If you choose **Specify at Deployment**, provide a name for the host in the **Host Name** field. If you choose **Auto-Generate**, specify a template for the name in the **Host Name Template** field.
 - e. If you are adding a node to a hyperconverged service, specify the **Host Name Selection** under **SVM OS Settings** and provide details about the hostname, as you did for the **OS Settings**.
 - f. In the **IP Source** box, provide an IP address.
 - g. Under **Hardware Settings**, in the **Node Source** box, select **Node Pool** or **Manual Entry**.

- h. In the **Node Pool** box, select the node pool. Alternatively, if you chose **Manual Entry**, select the specific node in the **Choose Node** box.
- i. Under **PowerFlex Settings**, specify the **Fault Set** for a node:
 - Select **PowerFlex Manager Selected Fault Set** to have PowerFlex Manager choose the fault set name based on the template settings.
 - Select ***fault-set-name*** to choose one of the fault sets in the service.

You can add nodes within a fault set, but PowerFlex Manager does not allow you to add a new fault set within the same service. To add a new fault set, deploy a separate service with settings for the fault set you want to create.

- j. Click **Next**.
- k. Review the **Summary** page and click **Finish**.

If the node you are adding has a different type of disk than the base deployment, PowerFlex Manager displays a banner at the top of the **Summary** page to inform you of the different disk types. You can still complete the node expansion. However, your service might have suboptimal performance.

Based on the component type, the required settings and properties are displayed automatically and can be edited as permitted for a node expansion.

5. Click **Save** to deploy the PowerFlex node.
When the job is complete, confirm that the node deployed successfully.

Managing the PowerFlex Manager virtual appliance

Virtual appliance management enables you to configure syslog or SNMP trap forwarding, edit proxy node or DHCP settings, configure the alert connector, and more.

Configure syslog forwarding

You can configure PowerFlex Manager to forward syslogs it receives from system components to a remote network management system. PowerFlex Manager provides authentication through the configuration settings you provide.

About this task

You can configure PowerFlex Manager to forward syslogs to up to five destination remote servers. You can set only one forwarding entry per remote server.

You can apply forwarding filters that are based on facility type and severity level. For example, you can configure PowerFlex Manager to forward all syslog messages to one remote server and then forward syslog messages of a given severity to a different remote server. The default is to forward syslog messages of all facilities and severity levels to the remote syslog server.

Prerequisites

- Ensure your system components are configured to send syslog messages to the PowerFlex Manager virtual appliance.
- Obtain the IP address of the hostname for the remote syslog server and the port where the server is accepting syslog messages.
- If sending only some syslog messages to a remote server, you must know the facility and severity of the log messages to forward.

Steps

1. On the menu bar, click **Settings > Virtual Appliance Management**.
2. On the **Virtual Appliance Management** page, in the **Syslog** section, click **Edit**.
3. For **Host**, enter the destination IP address of the remote server to which you want to forward syslogs.
4. Enter the destination **Port** where the remote server is accepting syslog messages.
5. Select the network **Protocol** used to transfer the syslog messages. The default is **TCP**.
6. Optionally enter the **Facility** and **Severity Level** to filter the syslogs that are forwarded. The default is to forward all.
7. Click **Save** to add the syslog forwarding destination.

The **Virtual Appliance Management** page displays the configured details as shown below:

```
Syslog Forwarding <destination-ip>(<Facility><Severity Level>)
```

SNMP communication

PowerFlex Manager can receive SNMP traps from system devices. It can forward SNMP traps to one or more remote network management systems.

PowerFlex Manager and the network management system use access credentials with different security levels to establish two-way communication. Review the access credentials that you need for each supported version of SNMP. Determine the security level for each access credential and whether the credential supports encryption.

To configure SNMP communication, you need the access credentials and trap targets for SNMP, as shown in the following table:

If adding...	You need to know...
SNMPv2	Community strings by which traps are received and forwarded.
SNMPv3	User and security settings

For SNMPv2 traps to send from a device to PowerFlex Manager, you must provide PowerFlex Manager with the community strings on which the devices are sending the traps. If during resource discovery, you selected to have PowerFlex Manager automatically configure iDRAC nodes to send alerts, you must enter the community string that is used in that credential.

For a network management system to receive SNMP traps from PowerFlex Manager, perform the following configuration outside of PowerFlex Manager:

- For SNMPv2 traps, you must provide the community strings to the network management system.
- For SNMPv3 traps, you must provide the PowerFlex Manager engine ID, user details, and security level to the network management system.

PowerFlex Manager supports different SNMP versions, depending on the communication path and function. The following table summarizes the functions and supported SNMP versions:

Function	SNMP version
PowerFlex Manager receives traps from all devices including iDRAC	v2
PowerFlex Manager receives traps from iDRAC devices only	v3
PowerFlex Manager forwards traps to the network management system	v2, v3

 **NOTE:** SNMPv1 is supported wherever SNMPv2 is supported.

SNMP trap forwarding configuration supports multiple forwarding destinations. If you provide more than one destination, all traps coming from all devices are forwarded to all configured destinations in the appropriate format.

PowerFlex Manager stores up to 5 GB of SNMP alerts. Once this threshold is exceeded, PowerFlex Manager automatically purges the oldest data to free up space.

Configure SNMP

Perform this procedure to configure SNMP communication to enable PowerFlex Manager to receive and forward SNMP traps. To configure SNMP, specify the access credentials for the SNMP version you are using and then add the remote server as a trap destination.

Steps

1. On the menu bar, click **Settings > Virtual Appliance Management**.
2. On the **Virtual Appliance Management** page, in the **SNMP Trap Forwarding** section, click **Edit**.
3. To configure trap forwarding as SNMPv2:
 - a. Click **+ Add community string**.
 - b. In the **Community String** box, provide the community string by which PowerFlex Manager receives traps from devices and by which it forwards traps to destinations.

You can add more than one community string. For example, do this if the community string by which PowerFlex Manager receives traps differs from the community string by which it forwards traps to a remote destination.

PowerFlex Manager adds the new community string to the list. The **Created By** column shows the user that created the community string. Any community strings associated with credentials are listed. For these community strings, the **Used By** column shows the credentials that use the community strings.

4. To configure trap forwarding as SNMPv3:
 - a. Click **+ Add User**.
 - b. Enter the **Username**, which identifies the ID where traps are forwarded on the network management system. The username must be at least 16 characters.
 - c. Select a **Security Level**:
 - **Minimal** - No additional information is required.
 - **Moderate** - Provide an **MD5 Authentication Password** that is at least eight characters.

- **Maximum** - Provide an **MD5 Authentication Password** and a **DES Privacy Password** that is at least eight characters.

Note the current engine ID (automatically populated), username, and security details. Provide this information to the remote network management system so it can receive traps from PowerFlex Manager.

You can add more than one user.

5. In the **Trap Forwarding** section, click **Add Trap Destination** to add the forwarding details.
 - a. In the **Target Address (IP)** box, enter the IP address of the network management system to which PowerFlex Manager forwards SNMP traps.
 - b. Provide the **Port** for the network management system destination.
 - c. Select the SNMP **Version** for which you are providing destination details.
 - d. In the **Community String/User** box, enter either the community string or username, depending on whether you are configuring an SNMPv2 or SNMPv3 destination. For SNMPv2, if there is more than one community string, select the appropriate community string for the particular trap destination. For SNMPv3, if there is more than one user defined, select the appropriate user for the particular trap destination.

6. Click **Save**.

The **Virtual Appliance Management** page displays the configured details as shown below:

```
Trap Forwarding <destination-ip>(SNMP v2 community string or SNMP v3 user)
```

Importing Active Directory users

Before importing Active Directory users to PowerFlex Manager, you must create at least one directory service using PowerFlex Manager. PowerFlex Manager roles are not automatically mapped to Active Directory user roles. You must assign an appropriate role to each imported user. After the import, these users can log in to PowerFlex Manager.

About this task

The format to log in is: **<Directory Service Name>/ <user name>**, and then type the password.

If an imported user is deleted from Active Directory, that user is not automatically deleted from PowerFlex Manager. The deleted user cannot log in to PowerFlex Manager, and you must remove the user manually from the user list. Importing an already imported user does not have any effect. The user role also remains the same.

Steps

1. On the menu bar, click **Settings** and then click **Users**.
2. On the **Users** tab, click **Import Active Directory Users**.
The **Import Active Directory Users** page is displayed.
3. Select a specific directory source from the **Directory Source** drop-down list to import the users from the selected directory source.
4. Select any of the following options from the **View** menu to filter the search results:
 - **All**—Displays both users and groups
 - **Users**—Displays only users
 - **Groups**—Displays only groups
5. Under the **Available Users/Groups** section, enter a user or group name in the **Find a User/Group** field to search for one or more users or groups in the selected directory.

To perform an explicit search for a particular user or group, type the exact letters of the user or group.

To search for a collection of users or groups, you can use a wildcard (*) character in your search. Wildcards are not automatically added to search strings. Therefore, you must enter them explicitly.

NOTE:

If possible, do not use a wildcard (*) at the start of a search string, as this search pattern can slow down the search significantly. For example, avoid a search string such as **'*smith'**. This type of search can be very slow, and, sometimes, may cause the search to timeout and return an error. Embedded or trailing wildcards are much faster. For example, you might improve search performance by entering **'john*'** or **'user*abc'**.

You can also speed up searches by specifying the lowest point in the AD tree as the base DN for the search.

6. Select the users or group you want to import and click the forward arrow (**>>**) button.
The selected users or groups are added to the **Users/Groups to be imported** section.
7. To assign a role to all the users or groups, select the users or groups and select any of the following roles from the **User Role** drop-down list: **Read Only**, **Administrator**, **Operator**, or **Standard**.
 - To apply specific roles, select the role from the **Role** menu beside the user or group name.
 - To view the imported group, select **All Groups** from the **Filter by Group** menu.

In a single import operation, if you import a user individually and as part of group, the role that is assigned to the user individually precedes the role that is assigned to the group.

Restarting the PowerFlex Manager virtual appliance

Use this task to restart PowerFlex Manager.

About this task

To restart the virtual appliance, you must be a user with the administrator role. The restart operation logs off all other users and cancel any running jobs.

Steps

1. On the menu bar, click **Settings**, and then click **Virtual Appliance Management**.
2. On the **Virtual Appliance Management** page, click **Reboot Virtual Appliance**. A message displays confirming that you want to restart the virtual appliance.
3. Click **Yes** to confirm. The system restarts.
4. Once the reboot is complete, click **Click to log in** and provide your credentials.

Troubleshooting and support

Identify problems, and resolve issues through generating troubleshooting bundles and accessing application logs. If you cannot resolve an issue, see [Support](#) to search the knowledge base for FAQs, Tech Alerts, and Tutorials.

Generate a troubleshooting bundle

A troubleshooting bundle is a compressed file that contains logging information for PowerFlex Manager managed components. If necessary, download the bundle and send it to Dell EMC Support for issue debugging.

About this task

The troubleshooting bundle includes the following logs:

- PowerFlex Manager application
- Alert connector
- PowerFlex Gateway
- iDRAC life cycle
- Dell EMC PowerSwitch switch
- Cisco Nexus switch
- VMware ESXi

If the alert connector is configured to send alerts through Secure Remote Services, the default location where the bundle is sent is Secure Remote Services. To configure the alert connector, go to the **Alert Connector** section on the **Virtual Appliance Management** page.

If you are not using Secure Remote Services, you can select the bundle to download locally.

Steps

1. On the menu bar, click **Settings > Virtual Appliance Management**.
2. On the **Virtual Appliance Management** page, click **Generate Troubleshooting Bundle**.
3. Select either **Send to Configured Secure Remote Services** or **Download Locally**.
4. If providing a local path use one of the following formats:

For...	Use this path...
CIFS	<code>\\IP Address\Any folder</code> For example: <code>\\192.168.1.1\uploadDirectory</code>
NFS	<code>IP Address:/Any folder</code> For example: <code>192.168.1.1:/var/nfs/uploadDirectory</code>

5. Click **Test Connection** to verify the connection to the CIFS or NFS share before generating the bundle.
6. Click **Generate**.

Access the application logs

PowerFlex Manager provides activity logs of user- and system-generated actions to use for troubleshooting. By default, log entries display in order of occurrence. You can access the logs through PowerFlex Manager or through the command line.

About this task

The logs are categorized in PowerFlex Manager as follows:

- Deployment
- Infrastructure or hardware configuration
- Infrastructure or hardware monitoring
- Licensing
- Miscellaneous
- Network configuration
- Template configuration
- Appliance configuration
- Security

Steps

1. To access logs in PowerFlex Manager, complete the following steps:
 - a. On the menu bar, click **Settings > Logs**.
 - b. Under **View Categories**, select the category of logs you want to access.
 - c. Click **Export All**.
The logs export to a CSV file.
 - d. View the **Severity, Category, Description, Date and Time, and User**.
2. To access PowerFlex Manager logs through the command line, complete the following steps:
 - a. Access the PowerFlex Manager virtual appliance.
 - b. Log in to PowerFlex Manager as **delladmin**.
 - c. Go to the logs you want, and copy them to a local directory in WinSCP.

Logs	Log location
<ul style="list-style-type: none"> • Resource discovery • Templates • Services 	/opt/Dell/ASM/logs/asmManager.log
<ul style="list-style-type: none"> • Service deployments • Failures found in <code>exception.out</code> files in the deployment directory for the service ID 	/opt/Dell/ASM/deployments/"serviceID"/deployment log
<ul style="list-style-type: none"> • Upgrade <ul style="list-style-type: none"> ◦ grep YUM alcm.log to search for upgrade errors • Back up and restore • Virtual appliance lifecycle management log (DHCP and or NTP) 	/opt/Dell/ASM/logs/alcm.log
<ul style="list-style-type: none"> • Exceptions not handled by other services • 500 errors 	/opt/Dell/ASM/logs/catalina.out
<ul style="list-style-type: none"> • Deployments • Firmware upgrades 	/opt/Dell/ASM/deployments
Health monitoring	/var/log/nagios
Alert connector	/opt/Dell/ASM/logs/alert_connector_rest_controller.log

Logs	Log location
	/opt/Dell/ASM/logs/ alert_connector_adapter.log

Exporting a configuration report for all resources and services

About this task

You can download a PDF report that lists configuration details for all resources and services.

The configuration report shows the result of various configuration checks that PowerFlex Manager performs against the system. You can use the report to troubleshoot your resources and services.

The report shows the following kinds of information:

Column name	Description
Resource Name	Provides the name of the resource for which the checks have been run. This column corresponds to the Resource Name shown on the Resources page.
Asset/Service Tag	Provides the asset or service tag of the resource for which the checks have been run. This column corresponds to the Asset/Service Tag shown on the Resources page.
Management IP	Specifies the Management IP of the resource for which the check is run. For a PowerFlex cluster, the management IP address is the MDM cluster IP address. For a PowerFlex Gateway, the management IP address is the PowerFlex Gateway IP address.
Result	Shows the result of the check (PASS or FAIL).
Severity	Indicates the severity of the check. The severity is based on the result. The severity levels are INFO, WARNING, HIGH, or CRITICAL. If the result is PASS, the severity is INFO. If the result is FAIL, the severity depends on the type of check. PowerFlex Manager supports only CRITICAL checks.
Details	Provides a description of the check that was run.
Affected Resources	Gives a list of the IP addresses or unique identifiers of resources that are impacted by the check. The list of affected resources helps with troubleshooting.

Steps

1. On the menu bar, click **Resources**.
2. Click **Export Report**.
3. Select **Export Configuration PDF Report** from the drop-down list.

The selected report downloads.