Dell EMC SmartFabric Services User Guide

Release 10.5.3

November 2021
Rev. A01

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.**

© 2020 - 2022 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Contents

1

About this guide

This guide provides information regarding the integration of SmartFabric Services (SFS) with Dell EMC VxRail, Dell EMC PowerEdge MX, Dell EMC Isilon/Dell EMC PowerScale devices, and other supported solutions. It covers the following details:

- SFS concepts and its components for leaf and spine deployment
- Description, configuration information, limitations, and restrictions of SFS for each solution
- Command reference for all the SFS commands

To use this guide, you must have a good knowledge of Layer 2 (L2) and Layer 3 (L3) networking technologies, and data center deployments.

This document may contain language that is not consistent with current guidelines of Dell Technologies. There are plans to update this document over subsequent releases to revise the language accordingly.

Text and Syntax Conventions

This guide uses the following conventions to describe text and command syntax.

Bold text	UI elements that you click or select
> (right angle bracket)	Hierarchy of menu selections
`Keyword`	Keywords are in Courier (a monospaced font) and must be entered in the CLI as listed
parameter	Parameters are in italics and require a number or word to be entered in the CLI
{X}	Keywords and parameters within braces must be entered in the CLI
[X]	Keywords and parameters within brackets are optional
x|y	Keywords and parameters separated by a bar require you to choose one option

Related Documents

Use the following documentation in addition to this guide to get complete information about the SmartFabric NVMe Controller capabilities:

Table 1. Related Documents

Related Documentation	Link
• *Dell EMC SmartFabric OS10 User Guide* • *Dell EMC SmartFabric OS10 Installation, Upgrade, and Downgrade Guide*	SmartFabric OS10 Documentation
Dell Technologies VxRail Documentation	Dell Technologies VxRail Networking Solutions
Networking Solutions Support Matrix	Support Matrix
PowerEdge MX Documentation	PowerEdge MX Manuals and Documents
PowerScale/Isilon Documentation	PowerScale OneFS Info Hub
OpenManage Network Integration Documentation	OMNI Documentation

Dell EMC Demo Center

The Dell EMC Demo Center is a scalable, cloud-based service that provides 24/7 self-service access to virtual labs, hardware labs, and interactive product simulations. Several interactive demos are available on the Demo Center. Contact your Dell account team to get access to the Demo Center.

Documentation Feedback

Dell Technologies strives to provide accurate and comprehensive documentation and welcomes your suggestions and comments. You can provide feedback in the following ways:

- Online feedback form—Rate the documentation or provide your feedback on any of product documentation pages at www.dell.com/support.
- Email—Send your feedback to networkingpub.feedback@dell.com. Include the document title, release number, chapter title, and section title of the text corresponding to the feedback.

To get answers to your questions related to Dell Networking Solutions through email, chat, or call, go to Dell Technologies Technical Support page.

Acronyms

The following acronyms are used throughout this guide:

Table 2. Acronyms

Acronym	Expansion
API	Application Programmable Interface
BGP	Border Gateway Protocol
CLI	Command-Line interface
GUI	Graphical User Interface
ICL	Interchassis link
ISL	Interswitch link
IOM	Input/Output Module
LACP	Link Aggregation Control Protocol
LAG	Link aggregation group
L2	Layer 2
L3	Layer 3
RPVST+	Rapid Per-VLAN Spanning-Tree Plus
SFS	SmartFabric Services
STP	Spanning-Tree Protocol
VLAN	Virtual LAN
VLTi	Virtual Link Trunking interconnect
VXLAN	Virtual extensible LAN
VTEP	VXLAN tunnel endpoint

2

Change history

The following table provides an overview of the changes to this guide from a previous release to the 10.5.3.0 release. For more information about the new features, see the respective sections.

Table 3. New in 10.5.3.0

Revision	Date	Feature	Description
A00	2021–10-12	Disable VxRail manager integration	SFS allows you to onboard VxRail nodes manually on to the fabric.
		Configure spine switch as a leader	SFS allows you to set the Preferred Leader flag for a spine switch in addition to the existing support for leaf switch.
		Enhanced SFS GUI	SFS GUI is enhanced to provide the following features: ● Enhanced SFS GUI design ● SFS summary view ● Fabric compliance information ● Switch management
		Dynamic discovery of storage devices	SFS dynamically discovers the storage devices and configures the discovered interfaces as untagged members of VLAN 4091.
		Configure IPv6 parameters for VXLAN networks	SFS allows you to configure L3 VXLAN network with IPv6 details.
		Configure BFD	SFS allows you to configure BFD as part of BGP configuration.
		Configure multisite fabric interconnect	SFS allows you to configure multisite fabric interconnect between the fabrics.

3

SFS fundamentals

This chapter provides information about the fundamentals of SFS including overview, supported topologies and platforms, network fabric formation, and its supported solutions.

SFS overview

SFS is a SmartFabric OS10 feature that provides network fabric automation and API-based programming capabilities. SFS has different personalities that can be integrated with systems including VxRail, PowerScale, generic PowerEdge servers, PowerStore, storage, and MX servers. SFS integrated with these solution-specific deployments delivers autonomous fabric deployment, expansion, and life cycle management.

SFS has two types: SFS for PowerEdge MX and SFS for leaf and spine. The following sections focus on concepts that are related to SFS for leaf and spine:

SFS for leaf and spine is supported on S-series and Z-series Dell EMC PowerSwitches. See Supported platforms for a complete list of supported platforms. SFS for leaf and spine has two personalities:

L2 single rack personality

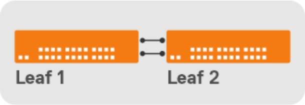

Single Rack Fabric

NOTE: This personality is not available for deployments after OS10.5.0.5 release. All single rack and multirack deployments from 10.5.0.5 release and later use the L3 personality.

SFS deployments on OS10 releases from 10.4.1.4 to OS10.5.0.5 support only L2 single rack.

- Provide fabric automation for a single pair of leaf switches.
- SFS deployment is limited to a single rack and cannot be expanded to a multirack deployment.

  NOTE: When you upgrade switches with this personality enabled, they operate in the L2 single rack personality only.
- SFS L2 single rack personality is enabled by running a Python script in the OS10 Linux shell. See Enable L2 personality for more information.

L3 multi rack personality

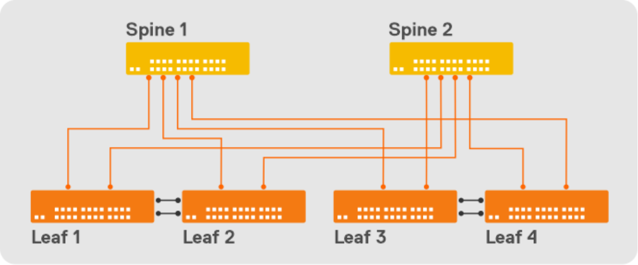

Multi Rack Fabric

All SFS deployments using OS10.5.0.5 and later releases support SFS L3 personality and the capabilities are as follows:

- Provides fabric automation for leaf and spine.
- Allows SFS deployment in a single rack and expand to multirack as required.
- Allows you to enable SFS using CLI, API, or GUI.

  NOTE: The support to enable SFS using GUI is not available from 10.5.3.0 release.

NOTE: In an MX-based deployment, the fabric is configured using OpenManage Enterprise-Modular (OME-M) UI, see *Dell EMC OpenManage Enterprise-Modular Editions User Guide* available in PowerEdge MX Documentation.

See the Solutions Support Matrix for a complete list of solutions that can be onboarded onto the fabric.

This guide covers the following SFS qualified solutions:

- SFS deployment with VxRail
- SFS deployment with PowerEdge MX
- SFS deployment with Isilon/PowerScale

SFS terminology

The terms master and leader are used interchangeably in this document. Leader is the same as master.

SFS GUI

OS10 has support for SFS GUI to set up the initial SFS configurations in a L3 leaf and spine topology. You access the SFS GUI using the latest version of the following browsers:

- Google Chrome
- Mozilla Firefox
- Microsoft Edge

For more information about SFS GUI, see Access fabric setup configuration.

Supported network topologies

Following are the supported network topologies for SFS L3 multirack deployment:

- One leaf switch pair without spine switch
- Multiple leaf switch pairs with a two or more spine switches

SFS supported platforms

SFS is supported on the selected S-series and Z-series PowerSwitches for leaf and spine deployments. The platform support varies depending on the solutions such as VxRail, PowerScale, and so on. The supported platforms are listed in *Supported Switches* section under respective solution chapters.

SFS and supported solutions

SFS has different levels of integration with the following qualified solutions. See the Solutions Support Matrix for the latest supported versions for all the qualified solutions.

Table 4. SFS and supported solutions

Qualified solutions	Dynamic discovery	Onboarding type
VxRail	Yes	Dynamic or Import from Fabric ⓘ **NOTE:** Use Import from Fabric option to manually onboard the VxRail nodes.
PowerStore X/T	Yes	Import from Fabric
PowerMax	Yes	Import from Fabric
Other third-party devices	Yes	● Import from Fabric, if NIC supports LLDP. ● Static, if NIC does not support LLDP.
PowerEdge MX	NA	NA

ⓘ **NOTE:** Any third-party device that supports standard Ethernet ports and protocols can be onboarded onto the solution using the onboarding procedure, see Onboard a server.

ⓘ **NOTE:** In PowerEdge MX, the servers are discovered and onboarded through OME-Modular.

Dynamic Discovery—Devices that support dynamic discovery send an LLDP TLV. Devices that do not send the LLDP TLV must be manually added to the fabric.

Onboarding—Onboarding is the process of adding devices to the fabric through the server interface profiles creation. Supported devices are automatically populated in the SFS GUI and OMNI by MAC address, switch, and switch port number for onboarding to the fabric. For VxRail, the SFS and VxRail Manager automates the onboarding process. You can also manually onboard the VxRail using **Import from Fabric** option. PowerStore systems support dynamic discovery and you can onboard the server using the **Import from Fabric** feature in SFS and OMNI, see Import SmartFabric discovered server interfaces. Hosts running ESXi may be onboarded using the Import from fabric option only if the hosts are already connected to vCenter. Other devices are manually onboarded by specifying the switch and switch port number for each interface, see Create server interface profile.

Creating a network fabric

When you enable SFS on the switches in leaf and spine architecture, a single network fabric is built with all the discovered switches using industry-standard L2 and L3 protocols. SFS supports the following automation capabilities in a leaf and spine topology:

- Elects one switch from the fabric as the leader switch. From 10.5.3.0 release, spine switch is also part of the SFS leader election and you can also set the spine switch as the Preferred Leader.
- Creates the infrastructure VLANs on the nodes.
- Allocates all the necessary internal IP addresses for leaf and spine configurations.
- Autoconfigures necessary BGP for all the relevant leaf and spine switches.
- Enables leaf and spine for underlay and overlay.
- Provides workload orchestration of server discovery and binding server profiles to networks.

See Internal components and entities for more information about infrastructure networks and fabrics that are created by SFS. Also, you can use SFS show commands to view and verify fabric-related configuration.

After the fabric is created, you can configure uplinks, jump host, and onboard servers to the fabric using the SFS GUI.

Server discovery and onboarding

SFS discovers and onboards a server based on the LLDP from the server when connected to the fabric.

Onboarding a server involves creating a server profile and assigning the networks (VLANs) to a specific switch port connected to the server NIC port.

You can onboard a server in the following ways:

- Dynamic onboarding
- Static onboarding

Dynamic server onboarding

When the servers are connected to the fabric, SFS discovers the servers automatically.

From 10.5.3.0 release, SFS dynamically discovers PowerMax devices using the LLDP TLV received from these devices and configures the discovered interfaces as the untagged member of the VLAN 4091—default client management network.

ⓘ **NOTE:** This feature is applicable only for L3 personality.

SFS discovers a host as a known server TLV in LLDPDUs sent through the connected ports. Following are the list of known servers discovered by SFS:

- VxRail

> **NOTE:** By default, SFS discovers and onboards the VxRail server automatically. From 10.5.3.0 release, SFS provides an option to onboard the VxRail servers manually by disabling the VxRail Manager Integration option. For more information, see Edit default fabric settings. This option is applicable only for SFS L3 personality.

- PowerStore X
- PowerStore T
- PowerMax

When a known server is discovered on the server-facing port, SFS applies the server profiles or server interface profile configurations.

Dynamic discovery of unknown servers

With OS10.5.2.2 and later releases, SFS dynamically discovers unknown servers using the standard LLDPDUs sent out through the connected ports. Unknown server is a host that does not send a Dell custom TLV in LLDP frame. Upon discovery, the client management network is configured by default.

> **NOTE:** If the LLDP TLVs are advertised with both bridge and routing capabilities, SFS considers that the interface is connected to a switch or router. Hence, the client management network, VLAN 4091, is not configured on that interface.

For dynamic discovery, the Port-id TLV in LLDP packet is mandatory. SFS matches the port-ID TLV value in the LLDP packet against the `Interface` field in the server interface profile to onboard the unknown server. After the server interface profile is configured, the onboard criteria for unknown server is same as onboarding a known server.

> **NOTE:** Ensure that the LLDP port-ID TLV value of the unknown server is unique to avoid misconfiguration.

- If both known and unknown LLDP neighbors exist on the same interface, SFS discovers these servers as known and unknown servers. During onboarding, if the server profile matches with both the known and unknown discovered servers on the same physical interface or with the same server interface profile, the order of onboarding is as follows:
  - Known discovered server
  - Unknown discovered server

  If you offboard the known server, SFS checks for the unknown server and onboards it if available.
- If both known and unknown LLDP neighbors exist with the same port ID, SFS discovers both the known and unknown servers. During the onboard process, the order of onboarding is as follows:
  - Known discovered server
  - Unknown discovered server

  If you offboard the known server, SFS checks for the unknown server and onboards the server if available.
- If there is a static onboarding server profile that is configured on an interface where the known and unknown LLDP neighbors exist, SFS discovers both the LLDP neighbors. During onboarding if there are conflicts on the physical interface or server interface profiles, the priority order of onboarding is as follows:
  - Static onboarding
  - Known discovered server
  - Unknown discovered server

  If you offboard the statically onboarded server, SFS checks the known servers list first and then the unknown server list to onboard the device based on the priority order.

Static server onboarding

SFS supports static onboarding of server on assigned ports instead of LLDP-based discovery. This option is used for onboarding servers that are not discovered by SFS.

To statically onboard a server, you must assign an interface of the leaf switch to which the server is onboarded. See *Onboard nondiscovered server interfaces* section in Onboard a Server for more information about static onboarding using GUI.

Configuration notes:

- STP is disabled on the server-connected ports.
- The bonding can be static or LACP.
- All existing bonding modes are supported on a statically onboarded server.

- VXLAN does not support STP on access ports and it is not applicable for L3 Fabric. For the VXLAN type of network, you cannot configure STP; the network topology must remain loop free.
- All existing network types are allowed to be onboarded on statically onboarded servers.
- Since onboarding is static, when server is moved there is no support for moving the configurations along with the server.
- The port-role for the statically onboarded server is `EndHost` or `GenericEndHost`.
- When the server profile or server interface profile is deleted, all the impacted interfaces are brought to default configuration.

Server profile and server interface profile

Server profile and server interface profiles are logical entities of SFS. You configure these entities when onboarding a server to the fabric.

Server interface profile represents the server NIC ports that are connected to the leaf switches. For a server interface profile, a server interface ID must be configured using the MAC address of the server NIC port. SFS uses the server interface ID to identify the server NIC port and configure the required networks on the server NIC ports.

Server profile is a list of server interface profiles with a common bonding type. Supported bonding types are `AutoDetect` and `LACP`. You can configure a server profile using SFS GUI, see Server onboarding.

Bonding technology and NIC bonding

Bonding technology or type is a property of server profile and the supported bonding types are `Autodetect` and `LACP`. NIC bonding is a server interface profile property that determines whether the server interface profile to be in LAG or normal.

The table describes the behavior of SFS depending on the option of Bonding type and NIC bonding selected when configuring the server interface profile:

Table 5. SFS behavior for bonding type and NIC bonding

Bonding technology or type	NIC bonding	SFS behavior
AutoDetect	Disable	The VLAN is added to the physical interface when LACP PDUs are not received. After the LACP PDUs are received, the VLAN is applied to LAG.
AutoDetect	Enable	The system waits for LACP PDUs and after the LACP PDUs are received, the VLAN is applied to LAG.
LACP	Disable	The VLAN is added to the physical interface.
LACP	Enable	The system waits for LACP PDUs and adds the VLAN only on the LAG after it receives the LACP PDUs.

By default, the bonding type setting is `AutoDetect` and NIC bonding is disabled. Dell Technologies recommends that when you configure the bonding type for a server profile as `LACP`, you must enable NIC bonding.

Fabric back up and restore

You can backup the SFS configuration to an external device, and restore the fabric to a good state using the backed-up configuration during a failure or error.

The backup functionality allows you to backup all the user configuration to a text file in JSON format. The REST endpoints are available that returns a text file containing all the user configuration that was done through the REST interfaces.

The restore functionality allows you to restore the fabric to a last known good configuration. The REST endpoint is available that enables you to stream the backed-up configuration. The restore action wipes the data completely and restores the configuration on the switches from the backup file. The restore functionality is followed by leadership switchover. For the successful restore operation to occur, the restore activity should be applied to the same set of switches in the fabric. Also, the SFS personality should be the same when the backup was performed.

The backup and restore endpoints are accessible only for the users with `sysadmin` role. The REST payload is encrypted using the SSL protocol.

You can also restore the configuration for a fabric from SFS UI, see Restore.

Table 6. HTTP Methods

Functionality	HTTP Method	URL	Payload
Backup	GET	https://*<ip-address>*/redfish/v1/Dnv/Backup	Output is a text as stream and Content-Type is application or octet-stream.
Restore	POST	https://*<ip-address>*/redfish/v1/Dnv/Restore	Content-Type: application or octet-stream and payload contains the backup file.

Setting up SFS

This chapter explains the workflow to setup SFS including initializing SFS and enabling it on leaf and spine switches.

Prerequisites

Ensure that the following are met before enabling SFS on the switches:

- Configure the Out-of-band (OOB) management network on the leaf switches. OOB management network enables connections to the SFS GUI. Dell EMC PowerSwitch S3048-ON or N3248TE-ON can function as an OOB management switch with the OS10 factory default configuration. For more information about OOB management topology and connection details, see the respective *Solution Deployment Guides*.
- Verify the SmartFabric OS10 version on the leaf and spine switches and update them with recommended version listed in the Support Matrix.

SFS configuration notes

This section lists important behaviors, considerations, and recommendations you must know before deploying SFS:

Recommendations

See Internal components and networks for more information about the SFS components and networks that are created by SFS.

Recommendations and notes regarding the SFS components are:

- Dell Technologies recommends that you do not change or disable the STP settings on switches in SmartFabric mode. Any change to the STP settings results in reboot of all the switches in the fabric and disabling STP settings cause loops.
- Dell Technologies recommends using uplinks from a leaf pair as a best practice. Uplinks from leaf switches to external switches can be L2 or L3. All uplinks from spine switches to external switches must be L3.
- Except for Isilon/PowerScale back-end deployment, breakout feature is not supported on the leaf switches for the ICL (VLTi) or ISL (leaf and spine) connections. Use the default speed cables of the leaf switch when creating the ISL or ICL connections in a L3 fabric.
    - ISL example—If the default speed on the leaf switch port is 100G and spine switch port is 400G, use 100G cable for ISL connection. Since autobreakout feature is enabled on the spine switch, the port speed of the spine switch is set to 100G automatically.
    - ICL example—If the default speed of the leaf switches is 100G, use 100G cable for the ICL connection.
- Ensure that the Preferred Leader flag is set for the leaf or spine switches in the existing fabric before expanding the fabric to prevent the configuration loss. Once a leader is elected, it initiates all applications to automatically build the network fabric, and leader virtual IP address is advertised for applications to automatically discover the fabric through inband networks.
- If you want to breakout the ports to configure uplinks or jump hosts, ensure that the breakout ports are configured on the switch before creating and configuring these entities. You can configure breakout ports using SFS GUI.
- Ensure that the default networks that are created by SFS do not conflict with any networks in the existing deployment. If so, change the default networks using the instructions that are provided in Edit fabric settings.
- When you want to reboot any of the border leaf switch having L3 uplink with eBGP peer configuration that is connected to the external switch, Dell Technologies recommends shutting down the interface of the external switch that is the connected to that leaf switch before rebooting. After reboot and the BGP session is established within the fabric, bring up the interface of the external switch.
- Dell technologies recommends not to flap the admin status of the server connected ports using the SFS UI, OMNI UI and REST API in order to avoid traffic loss on those ports after switch reboot.
- In SmartFabric mode, you can create VLAN using the `interface vlan` command through OS10 CLI, but you cannot delete the VLAN from the CLI. Therefore, Dell Technologies recommends you to use the SFS GUI to create, edit, or delete a VLAN.

Default settings

- SFS creates VLANs from 4000 to 4094 for internal use. Do not use these VLANs for general use.
- SFS creates 172.16.0.0/16 and 172.30.0.0/16 networks for the leaf and spine network configuration for the default domain ID 100.
- By default, autobreakout feature is enabled on the spine switches in SmartFabric mode. Autobreakout works only with DAC and AOC cables.
- In SFS, all network QoS priorities are set to Iron by default and it cannot be changed.

SFS behavior

SFS elects one switch from the fabric as a leader switch and designates the remaining switches as the backup switches. In the event of a leader reload or failure, a new leader is elected from the backup switches using the keepalive information. The switches that are configured as Preferred leader have a higher priority to become the leader switch. If none of the switches are configured as the preferred leader, any switch can become the leader.

When you expand the fabric, the newly added switches may come up and form a fabric among themselves, and elect a leader before they are connected to the existing fabric. When the new fabric merges with the existing fabric, SFS elects a new leader switch for the combined fabric. If one of the new switches becomes the leader, it may overwrite the configuration in the existing fabric.

Preferred Leader behavior

Starting from OS10.5.3.0 release, you can also manually set the Preferred Leader flag for each switch (leaf and spine) in the fabric using SFS GUI. If there are spine switches in the fabric, Dell Technologies recommends you to set the Preferred Leader flag for all spine switches in the fabric for optimal performance. For more information about how to set the Preferred Leader, see switches.

> ⓘ **NOTE:** Before upgrading to OS10.5.3.0 version, ensure that the Management IP address is configured on all spine switches in the fabric.

- Before OS10.5.3.0 version, the Preferred Leader flag is set automatically on all the leaf switches when you create an uplink using SFS GUI.
- From OS10.5.3.0 version:
  - If the Preferred Leader flag is not set in any of the switches in the fabric, SFS sets the Preferred Leader flag automatically on all the spine switches when you create an uplink in the SFS GUI. If there are no spine switches in the fabric, the Preferred Leader flag is automatically set on the leaf switches when an uplink is created.
  - If any switches are already set as Preferred Leader, SFS does not change the Preferred Leader flag status of the individual switches.

Spanning tree considerations

- To avoid loops, SFS does not allow you to configure the same network on multiple uplinks.
- SFS supports RPVST+ and MSTP. The default spanning tree mode in SFS is RPVST+. Once the fabric is created, you can change the STP mode using SFS GUI, see Edit default settings.

  > ⓘ **NOTE:** When you change the mode using GUI, the whole fabric goes through a reboot and the new mode is set to MSTP. The reboot action impacts the traffic flow in the cluster.

- By default, RPVST+ is enabled on the uplink interfaces for L3 VLAN networks. When RPVST+ is enabled on the uplink, the total number of Port VLANs (PV) supported is 400.

BFD support

From OS10.5.3.0, SFS provides an option enable or disable Bidirectional Forwarding Detection (BFD) for external BGP peer routes. You can enable or disable BFD for a eBGP peer route policy when configuring uplinks, server interface profiles, and routing profiles. By default, BFD is disabled.

> ⓘ **NOTE:** You cannot modify the BFD timers in SFS:
>
> - For the S4100-ON series platform, the BFD interval is set to 500 ms with multiplier of 3 by default.

- For other platform switches, the BFD interval is set to 300 ms with a multiplier of 3 by default.

If you want to edit the BFD configuration for a routing policy, you must delete the existing routing policy and create a policy with the respective BFD configuration.

Configure SFS

By default, PowerSwitches boot in Full Switch mode when you power them up with SmartFabric OS10. This information explains how to setup a fabric.

1. Enable SFS and set roles on the switches.
2. Connect to SFS GUI and complete the relevant fabric configurations to setup the fabric. See SFS GUI for more information regarding login procedure and fabric configurations.

For SFS with Isilon/PowerScale deployment, you only have to enable SFS with roles on all the switches using CLI. All other SFS initial deployment operations using UI are not required. For more information regarding PowerScale deployments, see the related documents available in PowerScale Info Hubs.

Enable SFS on PowerSwitches

You can enable SFS on PowerSwitches through CLI, GUI, or REST API from OS10.5.0.5 and later versions. This option is applicable only for SFS L3 leaf and spine personality.

NOTE: The support to enable SFS using GUI is not available from 10.5.3.0 release.

After you enable SFS on all the switches in a leaf and spine deployment, a network fabric is created automatically with the default fabric settings. See Internal fabric component and networks for more information about the default settings.

Enable SFS using CLI

To enable SFS on a switch using CLI, run the `smartfabric l3fabric enable` command and set a role. For more information, see `smartfabric l3fabric enable`.

After you enable SFS on the switches and set a role, the system prompts for a confirmation to reload and boots in the SmartFabric mode. Enter `Yes` to continue. In SmartFabric mode, the CLI is restricted to global switch management features and monitoring. Using this command, enable SFS on all the switches with the corresponding role to create a fabric.

The following example shows how you can enable SFS on switches and set roles:

- Enable SFS on a switch with a spine role in CONFIGURATION mode.

  ```
  OS10(config)# smartfabric l3fabric enable role SPINE

  Reboot to change the personality? [yes/no]: yes
  ```

- Enable SFS on a switch with a leaf role in CONFIGURATION mode.

  ```
  OS10(config)# smartfabric l3fabric enable role LEAF vlti ethernet 1/1/4-1/1/5

  Reboot to change the personality? [yes/no]: yes
  ```

  In SFS, the two leaf switches are automatically configured as a VLT pair. The Ethernet interfaces 1/1/4 and 1/1/5 are the VLTi interfaces.

  If the VLTi interfaces are not contiguous, use the following syntax to enable VLT:

  ```
  smartfabric l3fabric enable role LEAF vlti ethernet 1/1/25,1/1/29
  ```

- (Optional) Enter the domain ID for the SFS fabric in CONFIGURATION mode.

  ```
  OS10(config)# smartfabric l3fabric enable role LEAF vlti ethernet 1/1/4-1/1/5 domain
  101

  OS10(config)# smartfabric l3fabric enable role SPINE domain 101
  ```

By default, domain ID is set to 100. You can edit the domain ID from SFS GUI. See Edit default fabric settings.

In Isilon or PowerScale back-end deployments, when enabling SFS in a leaf and spine topology, no VLTi configuration is required for the leaf switch. To set the leaf role for a switch, use the `smartfabric l3fabric enable role LEAF` command without the VLTi parameters.

Related Videos

Enable SFS on PowerSwitches using a CLI command:

Enable SFS using a CLI command

Enable SFS using RESTCONF API

You can enable SFS on OS10 switches using the RESTCONF API. For more information regarding general RESTCONF API operations, see *Dell EMC SmartFabric OS10 User Guide*.

Description

Enables SFS on the switches with leaf and spine roles.

RESTCONF endpoint

`/restconf/data/dell-smart-fabric:config-personality`

JSON content (spine)

```
{
"dell-smart-fabric:config-personality": {
"service-enable":true,
"role": "SPINE",
}
}
```

Example

```
curl -k -u admin:admin -H 'Content-Type: application/json' -i -X POST
-d \
'{"dell-smart-fabric:config-personality":
{"service-enable":true,
"role":"SPINE"}
}'
https://100.104.26.104/restconf/data/dell-smart-fabric:config-personality
```

The following example shows how to enable SFS on a leaf switch with an IP address (100.104.26.104) and ICL interfaces as 1/1/5 and 1/1/6.

JSON content (leaf)

```
{
"dell-smart-fabric:config-personality": {
"service-enable":true,
"mode": "L3 Fabric",
"role": "LEAF",
"icl": [
"ethernet1/1/5",
"ethernet1/1/6"
]
}
}
```

Example

```
curl -k -u admin:admin -H 'Content-Type:
application/json' -i -X POST -d \
'{"dell-smart-fabric:config-personality":{"service-
enable":true,"role":"LEAF",
"icl": [
"ethernet1/1/5",
"ethernet1/1/6"
]
}
}'
https://100.104.26.104/restconf/data/dell-smart-fabric:config-personality
```

Example—Usage of curl command on a linux server is as shown:

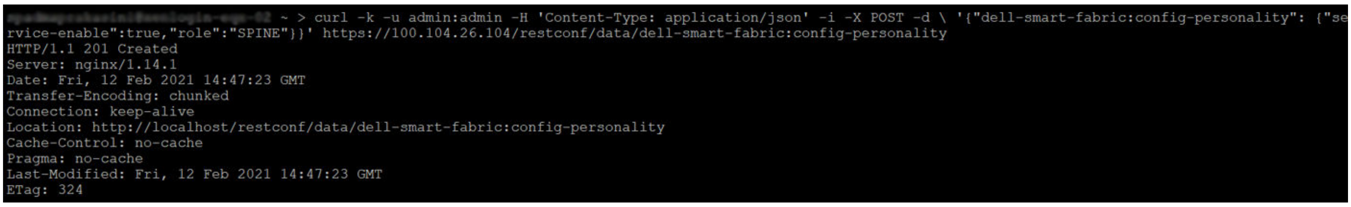

```
[redacted] ~ > curl -k -u admin:admin -H 'Content-Type: application/json' -i -X POST -d \ '{"dell-smart-fabric:config-personality": {"se
rvice-enable":true,"role":"SPINE"}}' https://100.104.26.104/restconf/data/dell-smart-fabric:config-personality
HTTP/1.1 201 Created
Server: nginx/1.14.1
Date: Fri, 12 Feb 2021 14:47:23 GMT
Transfer-Encoding: chunked
Connection: keep-alive
Location: http://localhost/restconf/data/dell-smart-fabric:config-personality
Cache-Control: no-cache
Pragma: no-cache
Last-Modified: Fri, 12 Feb 2021 14:47:23 GMT
ETag: 324
```

Enable SmartFabric Services L2 personality using script

In SFS L2 personality, you can enable SFS only using API. All deployments from SmartFabric OS10.4.1.4 to OS10.5.0.x version support only the single rack network fabric.

NOTE: The L2 personality is not available for new deployments after 10.5.0.5 release.

A python script is used to enable SFS. For more information about L2 personality commands and information, see *VMware Integration for VxRail Fabric Automation SmartFabric User Guide, Release 1.1* available in OMNI Documentation page.

Enable SFS using GUI

Starting from OS10.5.3.0 release, the option to enable SFS mode using GUI is not available. You can use the GUI to enable SFS in versions prior to 10.5.3.0.

To enable SFS using the GUI, following the instructions:

1. Enable RESTCONF API on the switch using the OS10 CLI.

   To enable RESTCONF API, use `rest api restconf` command in CONFIGURATION mode. See *Dell EMC SmartFabric OS10 User Guide* for more information about RESTCONF API.

2. Open a browser session, go to **https://*switch-mgmt-ip-address***.
3. Log in to a switch using the credentials that are created to access an OS10 switch. The default username and password is `admin`.
4. Enable SFS on the switch using the **Edit** option that appears in the upper-right side of the page.

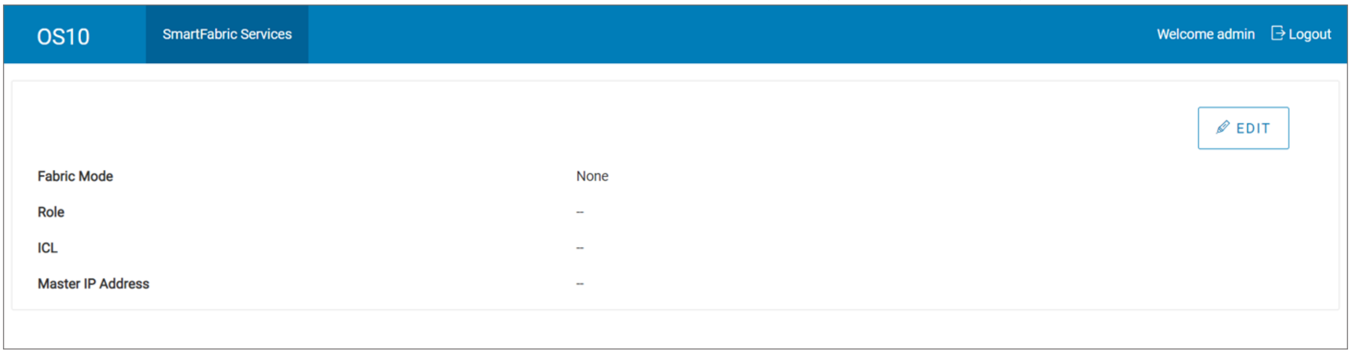

Fabric Mode None
Role --
ICL --
Master IP Address --

5. Enter the role of the switch and click **OK** to enable SFS.

Spine:

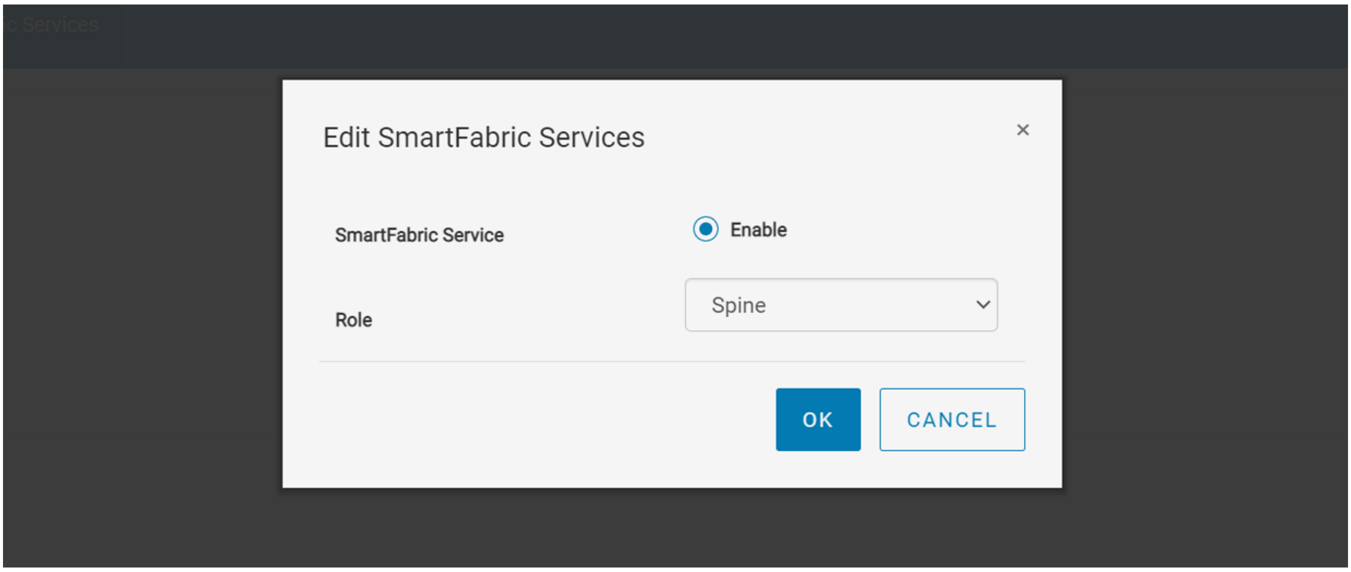

Edit SmartFabric Services

SmartFabric Service: ● Enable

Role: Spine

Leaf:

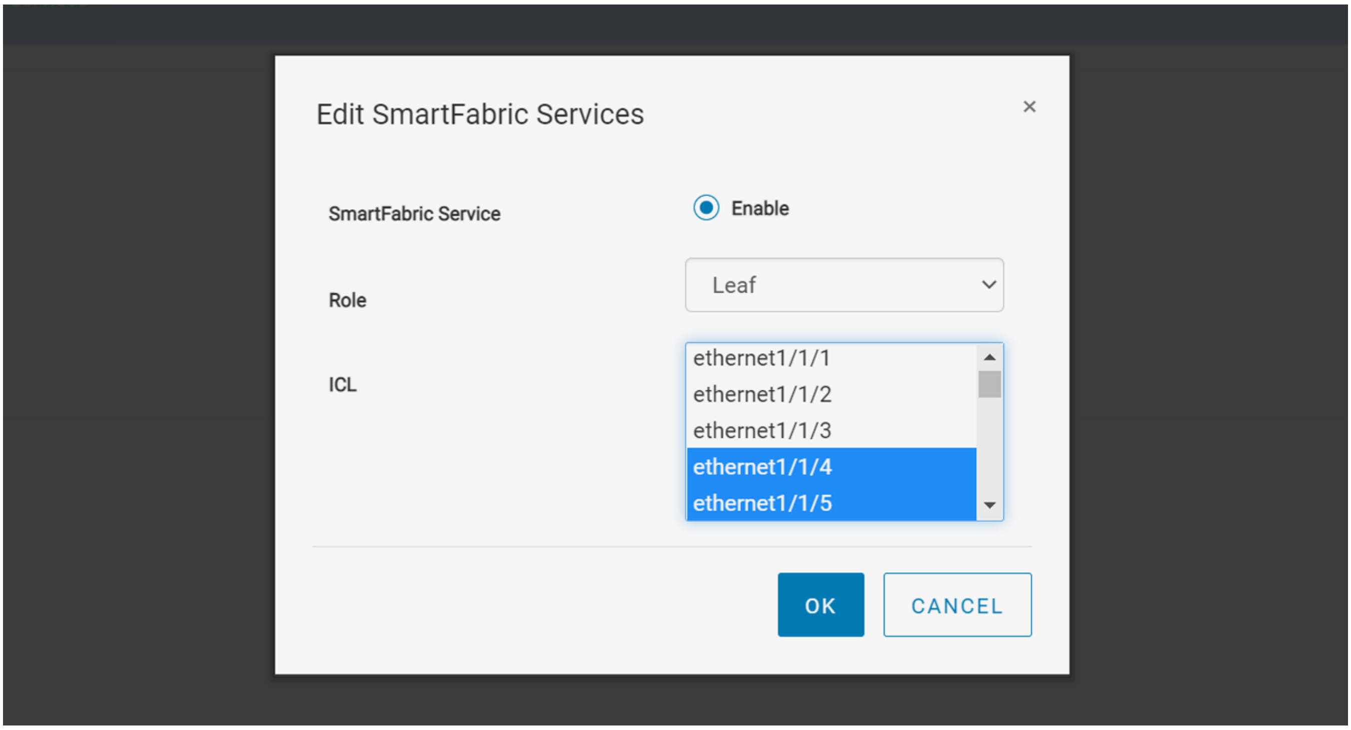

Edit SmartFabric Services

SmartFabric Service: ● Enable

Role: Leaf

ICL:
- ethernet1/1/1
- ethernet1/1/2
- ethernet1/1/3
- ethernet1/1/4 (selected)
- ethernet1/1/5 (selected)

After you enable SFS on a switch, the system reloads to apply the configuration.

6. Repeat the steps 1 to 4 on all the remaining switches to enable SFS and set role. All switches reload and forms a fabric.

Verify the switch operating mode

To verify that the switches are in SmartFabric mode, run the `show switch-operating-mode` command on each switch.

```
OS10# show switch-operating-mode
Switch-Operating-Mode : SmartFabric Mode
```

For more information regarding the CLI, see the show switch-operating-mode command.

Disable SFS using CLI

To delete the existing switch configuration and go back to Full Switch mode, run the `no smartfabric l3fabric` command on each switch.

The `no smartfabric l3fabric` command disables the L3 fabric personality. After you disable the L3 fabric in the switch, the system prompts for confirmation and reboots in Full Switch mode.

```
OS10(config)# no smartfabric l3fabric

Reboot to change the personality? [yes/no]: yes
```

Complete the fabric setup

After enabling SFS, you must configure logical entities to complete the SFS setup.

1. Log in to SFS GUI. For more information, see Deploy and manage a fabric.
2. Configure unique names for SFS components. For more information, see Update fabric names and descriptions.
3. (Optional) Configure breakout switch ports based on the deployment requirements. For more information, see Breakout ports.
4. Configure uplinks for connectivity. For more information, see Create uplinks.
5. Configure Jump host for client management based on the requirement. For more information, see Configure Jump host.
6. Onboarding a server by applying server interface profile. For more information, see Server onboarding.

SFS GUI also provides additional features that can be configured if needed. For more information, see Access fabric setup options.

5

Deploying and managing a fabric

You can use SFS GUI to set up a fabric in a leaf and spine topology. The SFS GUI helps you with the SFS deployment operations and management of the switches in a fabric.

Access fabric setup options

Follow the instructions to access fabric setup options:

You can launch the SFS GUI directly using the leader IP address. To get the IP address of the leader switch in the L3 fabric, use the `show smartfabric cluster` command. For more information, see SFS commands.

If you logged in to a switch in a fabric that is not the leader, the page displays only the fabric mode of switch and the IP address of the leader along with the link to the leader switch. Click the link to log in to the leader switch. All fabric configuration can be performed only from the SFS leader switch.

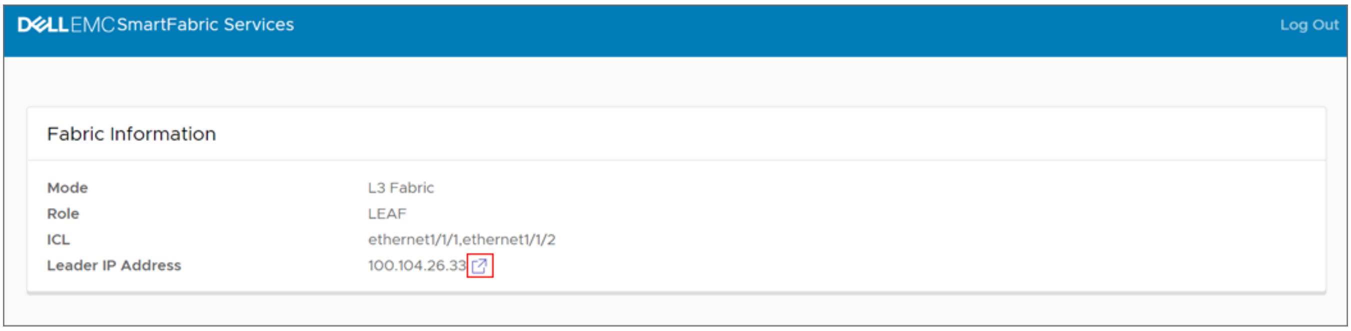

Fabric Information

Mode	L3 Fabric
Role	LEAF
ICL	ethernet1/1/1,ethernet1/1/2
Leader IP Address	100.104.26.33

1. Log in to SmartFabric Services GUI using the management IP address of any switch.
2. Enter the credentials that are created to access an OS10 switch from the console or through a network connection. The default username and password is `admin`. You can also use any user-configured accounts as credentials.

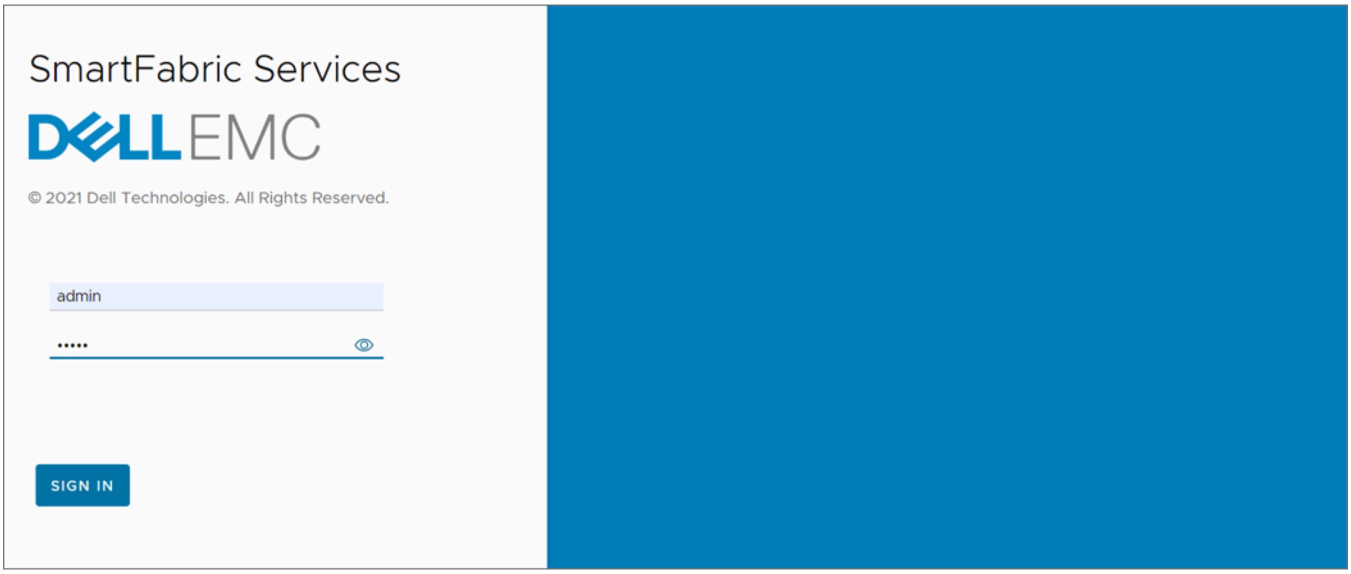

admin

•••••

Click **Logout** button at the upper right of the GUI to terminate the login session.

Upon successful login, the GUI displays the summary information.

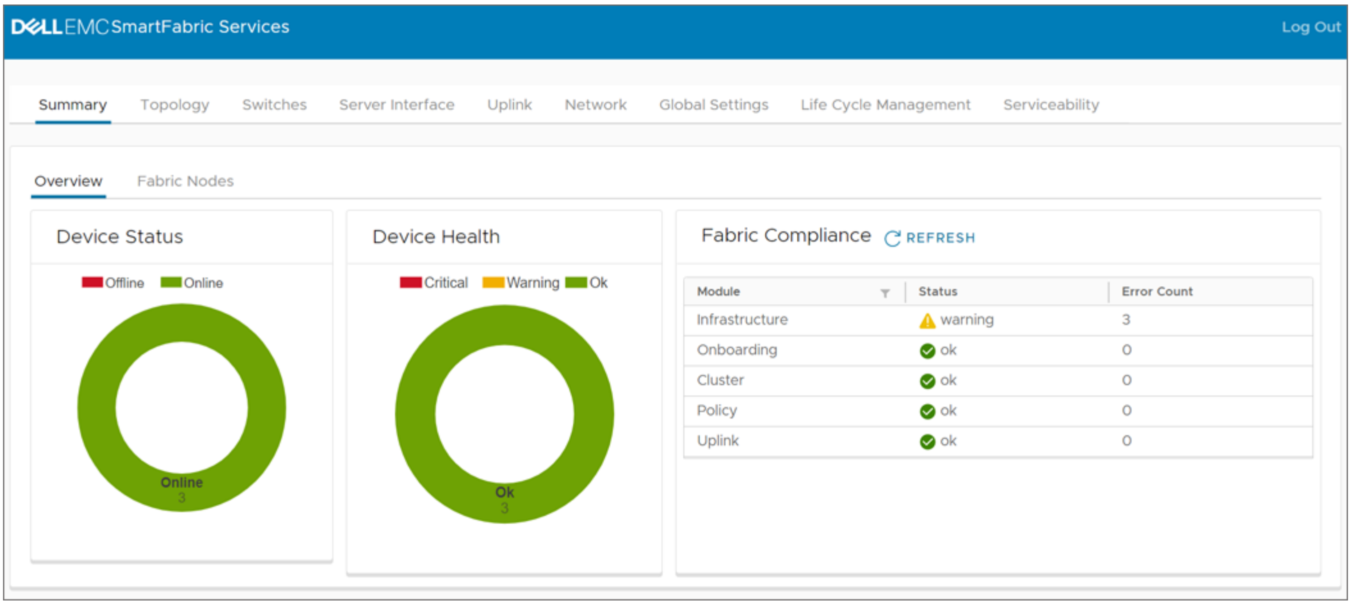

Module	Status	Error Count
Infrastructure	⚠ warning	3
Onboarding	✓ ok	0
Cluster	✓ ok	0
Policy	✓ ok	0
Uplink	✓ ok	0

SmartFabric overview

Starting from 10.5.3.0 release, SFS GUI displays a consolidated view of key metrics such as device status, device health, and the latest fabric compliance errors for the fabric. The **Overview** dashboard displays the following metrics:

Device Status—Displays the status of the all the switches that are deployed in the SmartFabric instances along with the number of switches in each status.

- Green—Indicates that the device is online.
- Red—Indicates that the device is offline.

Device Health—Displays the overall health of all the devices in the fabric.

Fabric Compliance—Displays the misconfiguration and compliance violations identified in the SmartFabric instance.

- Module—Name of the module in which the misconfiguration or compliance errors occurred.
- Status—Compliance status of each module in the fabric.
- Error Count—Number of errors in each module.

You can view the detailed list of all compliance errors in the SmartFabric instance from **Serviceability** page.

View node details

From SFS Summary page, you can view the switch details in the fabric.

From **Summary** > **Fabric Nodes**, you can view the list of spine and leaf nodes that are deployed in the fabric.

Click the Fabric IDs to view the detailed information of the spine and leaf switches connected in the fabric. The details include:

- Switch name
- Switch model
- OS10 version
- Switch role
- IP address of the switch
- Status of the switch (online or offline)

Click **Domain** at any time to update the fabric details.

View fabric topology

The **Topology** tab displays the graphical topology of the network fabric for the selected SmartFabric instance. You can also view the details of the switch in the fabric.

The L3 leaf and spine topology view that is created after you enable SFS. The topology view displays the switch icons with the hostname and the service tag information under each node and the link connectivity between the switches. Mouse over a fabric to see the detailed information about the leaf and spine switches, and the link connectivity.

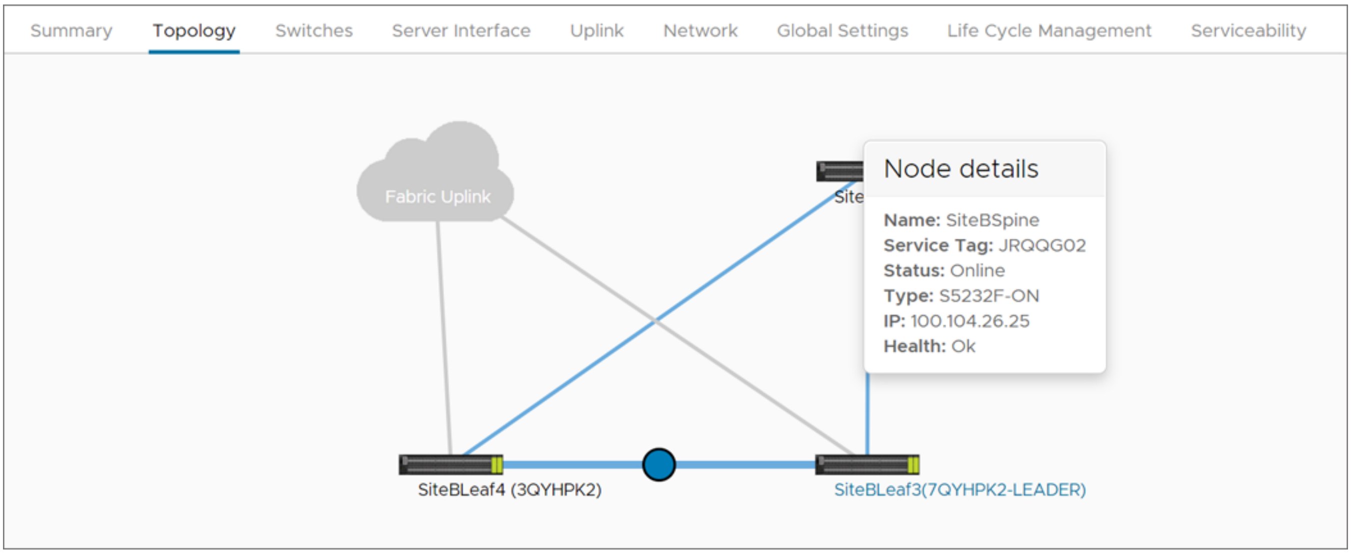

Node details

Name: SiteBSpine
Service Tag: JRQQG02
Status: Online
Type: S5232F-ON
IP: 100.104.26.25
Health: Ok

Manage switches in a fabric

You can manage the spine and leaf switches available in a fabric.

Set Preferred Leader

Dell Technologies recommends you to set a spine switch as the Preferred Leader for optimal performance. You can identify the switch that is set as the leader switch from SFS GUI. You can set the leaf or spine switches as a Preferred Leader from SFS GUI:

1. Click **Switches**.
2. Click the toggle switch to set the Preferred Leader for the leaf or spine switch. The system prompts for confirmation.

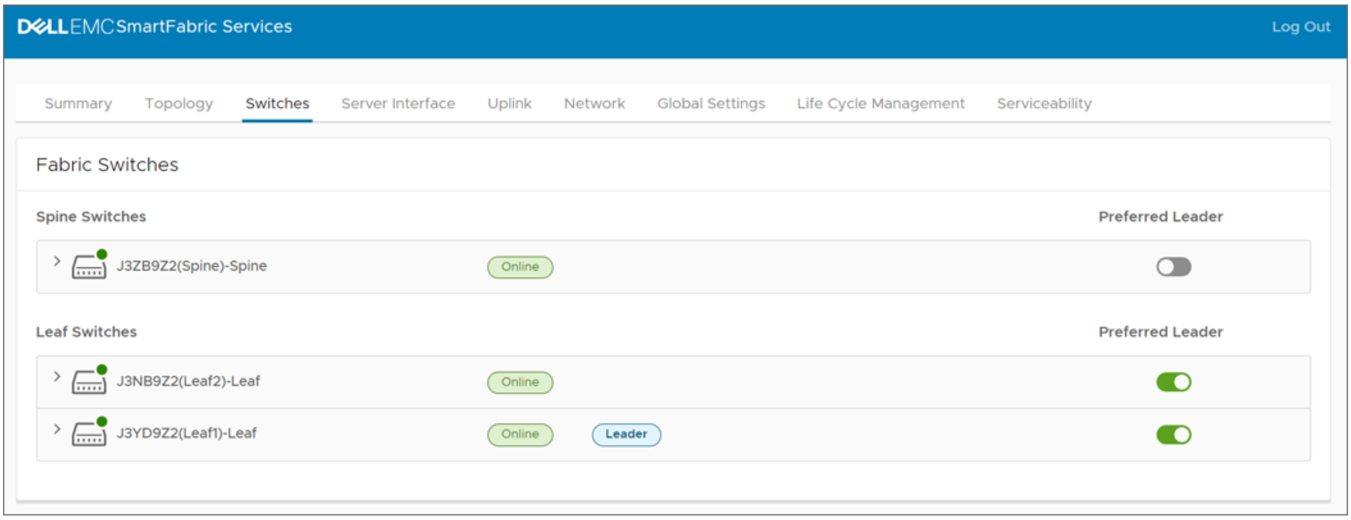

Fabric Switches

Spine Switches			Preferred Leader
J3ZB9Z2(Spine)-Spine	Online		○

Leaf Switches			Preferred Leader
J3NB9Z2(Leaf2)-Leaf	Online		●
J3YD9Z2(Leaf1)-Leaf	Online	Leader	●

3. Click **Ok** to confirm.

You can disable the Preferred Leader for any switch using the toggle switch. After confirmation, the Preferred Leader is disabled.

Change port configuration on a switch

Change the port configuration such as auto negotiation or MTU of a port or multiple ports on a leaf or spine switch:

△ CAUTION: **Changing the interface configurations can potentially cause a disruption in service. Ensure that you are aware of the network settings and the remote-peers connected to the interfaces before changing the MTU, auto negotiation, admin status. If the configuration does not match the connected peer switch, it can lead to connectivity issues.**

Edit a port configuration

1. Select **Switches**.
2. Select the spine or leaf switch by clicking the arrow to view more information.
3. Select a category (**All Ports** or **Unused Ports**).
4. Select a port from the category.
5. Click **Edit**.
6. Edit the following details:
   - (Optional) Name
   - (Optional) Description
   - Admin status
   - MTU
7. Click **Edit**.

Configure auto negotiation status

You can enable or disable the auto negotiation on a single port or multiple ports. Auto negotiation option is not applicable for port channel interfaces. If the selected list has a port channel interface when configuring auto negotiation, the system displays a warning to clear the port channel interfaces from the selected list. Clear the port channel interface selection and try again.

To enable auto negotiation:

1. From **All Ports**, select a port or multiple ports.
2. Click **Enable Auto Neg**. The system displays a warning message.
3. Click **Yes** to confirm.

   The system displays the stage-wise progress of the interface status.

To disable auto negotiation:

1. From **All Ports**, select a port or multiple ports.
2. Click **Disable Auto Neg**. The system displays a warning message.
3. Click **Yes** to confirm.

   The system displays the stage-wise progress of the interface status.

Set MTU value

Set maximum transmitting unit (MTU) for the port:

1. Select a port or multiple ports and click **Set MTU**.
2. Enter the MTU value and click **Set**. By default, the SFS value is set to 9216.
3. Click **Yes** to confirm.

   The system displays the action success or failure message.

Enable switch ports

You can view and manage the unused ports in the switches. To enable or disable unused ports:

1. From **Switches** tab, select **Unused Ports** category to view the list of unused ports available in the leaf or spine switch.
2. Select a port or multiple ports from the list.

3. Click **Enable Admin Status**.

   To disable the ports, select a port or multiple ports, and click **Disable Admin Status**. The system displays the change status and update success message on completion.

Dell Technologies recommends to:

- Enable the port status to operationally up before adding any devices to the port, if the port is disabled using the GUI.

  NOTE: Devices that are connected to the disabled port are not discovered.
- Ensure that the ports are UP before adding any switches, when you expand the leaf and spine fabric deployments.
- Ensure that the switch port is in UP, when onboarding a server to a leaf switch.

Breakout switch ports

You can configure breakouts for the Ethernet ports or port-group only on the leaf switches to connect to the external device or jump host. Use the following procedure to breakout switch ports:

NOTE: By default, the auto breakout feature is enabled in SmartFabric spine switches. SFS GUI does not provide an option to breakout ports in spine switches.

1. From **Switches** tab, click **Leaf Switches**.
2. Select a leaf switch from the list for which you want to breakout.
3. Click **Breakout Port & Jump Port** category.
4. Select a port that you want to breakout and click **Breakout Port**.

   NOTE: The existing configuration of the port is reset to default when you configure a breakout port.
5. Select the appropriate breakout option for the port from the list.
6. Click **Submit**. To view the details of the breakout ports, select a port to view the properties of the port.

   The system displays breakout port configured successful or failure message.

Configure jump host

A jump host is a designated port to which an external device such as a laptop can be connected. You can configure only one port in a leaf switch port as a jump host for the external device to connect to L3 fabric. Select any available port that is not part of an uplink, ICL, and port connected to a server in fabric.

In VxRail deployment, a jump host is primarily used to bring up a VxRail cluster. By default, all VxRail nodes are placed in the client control and client management networks. In VxRail deployment, the jump host port is placed in the client management network in order to reach the default VxRail Manager VM.

Use the following procedure to configure jump host:

1. From Leaf Switches, select the leaf switch from the list.
2. Select the **Breakout Ports & Jump Port** category.
3. Select the switch to view the properties and click **Jump Port**.
4. Enter the name for the jump host.
5. Select an interface of the leaf switch as the jump host.
6. Associate an untagged network with the jump host and click **Ok**.

Delete jump port

1. Select the leaf switch for which you want to delete the configured jump port.
2. Select the Jump port and click **Delete**.

   The system displays jump port deletion success message.

Configure server interface profile

See Server discovery and onboarding for more information about server onboarding and its types. You can onboard a server statically or dynamically.

Onboard a server onto the fabric

For more information about server onboarding types, see Server discovery and onboarding. You can onboard a server statically or dynamically.

Onboard discovered interfaces

Use the following procedure to dynamically onboard a discovered server:

1. From **Server Interface** tab, click **Create**.
2. Enter the server interface ID.

   ⓘ **NOTE:** The server interface ID must be unique. Dell Technologies recommends using the MAC address of the onboarded server interface without ":" .
3. Select the server profile. The template changes based on the server profile selection.
   - Existing server profile—Select the server profile ID from the list.
   - New server profile—Create a profile ID. Enter the unique string for the server profile ID and the bonding type.
4. Select one or multiple networks for tagged and untagged networks.
5. Select **No** for static onboarding.
6. Select Enable or Disable for NIC bonding.
7. Click **Create**.

Onboard nondiscovered server interfaces

Use the following procedure to onboard the server statically for nondiscovered server interfaces:

1. From **Server Interface** tab, click **Create**.
2. Enter the server interface ID.

   ⓘ **NOTE:** The server interface ID must be unique. Dell Technologies recommends using the MAC address of the onboarded server interface without ":" .
3. Select the server profile. The template changes based on the server profile selection.
   - Existing server profile—Select the server profile ID from the list.
   - New server profile—Create a profile ID. Enter the unique string for the server profile ID and the bonding type.
4. Select one or multiple networks that needs be tagged and untagged.
5. Select Enable or Disable for NIC bonding.
6. Select **Yes** for static onboarding.
7. Select the leaf switch and the interface on which the server is connected.
8. Select the routing protocol. The template changes based on the routing protocol.
   - None
   - eBGP—Enter the routing policy name, peer interface IP address, peer ASN, description (optional), and BFD configuration.
   - Static Route—Enter the routing policy name, network address, prefix length, next hop IP address, and description (optional).
9. Click **Create**.

Import SmartFabric discovered server interfaces

Automate onboarding of server interface profile by importing profiles that are discovered by SFS. From 10.5.3.0 release, this option **Import from Fabric** is available from SFS GUI to import and onboard the discovered server interfaces to the fabric.

When known servers are connected to the fabric, SFS discovers the servers automatically and onboards the discovered servers as part of this workflow. SFS discovers the hosts or servers as known using the originator field in the Dell custom LLDP TLV sent by the servers. From OS10.5.2.2 release, SFS discovers unknown servers and you can onboard the unknown servers using the **Import from Fabric** option. Onboarding unknown servers is applicable for the SFS L3 leaf and spine personality. Use this feature to onboard a known or unknown server.

- **Known server**—A known server is a host that sends a valid originator in Dell-specific (custom) TLV in LLDP frame that is recognized by SFS.
- **Unknown server**—An unknown server is a host that sends LLDP frames that do not include the Dell-specific TLV.

1. Click **Server Interface**.

2. Click **Import from Fabric**. **Discovered Server Interface** window appears with the list of discovered interfaces.

   NOTE: The interface that is already associated with a server interface profile is not listed in the discovery table.

3. (Optional) Select **Edit** icon at the end of each row to edit the server profile information of each interface. You can edit the following:
   - NIC bonding configuration
   - Static onboarding configuration
4. (Optional) Add server interface networks for the discovered server interface profiles if needed. For more information about adding networks, see *Add networks* section.

   NOTE: This action overwrites the existing networks of all the server interface profiles.

5. (Optional) Add the server profile for the interface if needed. For more information about adding server profile, see *Add to Server Profile* section.
6. Select one or multiple discovered interfaces and click **Update**.

Add interfaces to Server Profile

To add the discovered interfaces to a new or existing server profile:

1. Select one or more discovered interfaces, and click **Add to Server Profile**.
2. Select the server profile to which you want to add the discovered server interfaces.
   - Select **Existing Server Profile**—Select the **Server Profile Id** to associate the interface with the existing server profile, and click **Associate**.
   - Select **New Server Profile**—Enter the **Server Profile Id** and **Bonding Type** to associate the interface with the new server profile, and click **Associate**.
3. The system displays the server interface profile association success message.

Add Networks

To add the networks to the discovered interfaces:

1. Select one or more interfaces from the list, and click **Add Networks**.
2. Associate the networks with the discovered interfaces, and click **Add**.
   - Select one or multiple networks for **Tagged Networks**.
   - Select a single network for **Untagged Network**.

   NOTE: Add networks overwrite the existing networks of all the server interface profiles.

3. The system displays the server interface networks addition success message.

Remove from server profile

To remove the interface from the server profile, select one or more interfaces from the list, and click **Remove from Server Profile**.

Edit networks and ports in a server interface profile

You can edit the network and port configuration in a server interface profile. You can also view the detailed information of a server interface profile.

Select a server interface ID to view the properties of the profile on the right.

Edit networks on a server interface profile

1. Select the server interface ID from the list to view the detailed information.
2. Click **Edit Networks**.
3. Edit the **Untagged Network** and the **Network** configuration for the profile.
4. Click **Edit**.

   The system displays the server interface profile update success message.

Edit ports on a server interface profile

1. Select the server interface ID from the list, and click **Edit Ports**.
2. Edit the **Static Onboarding Option** and the **NIC Bonding** configuration for the profile.
3. Click **Edit**.

   The system displays the server interface profile update success message.

Configure and manage uplinks

Uplinks enable the network fabric to communicate with the external network. Before creating an uplink, ensure that the external network is configured with the L2 or L3 setup. Any ports available on the leaf switches may be used as uplinks, provided they are compatible with the corresponding ports on the external switches.

SFS supports eBGP and static routes profiles and LACP and static for uplink bonding. For more information, see Uplinks section.

Configure L2 Uplink

Follow the instructions to create a L2 uplink:

1. Click **Uplink** > **Create**.
2. Enter the uplink details:
   - Select the uplink port type as **L2**.
   - Enter the name and description (optional) for the L2 uplink.
3. Click **Next**.
4. Configure the port:
   - Select a rack in which you want to create the uplink.
   - Select one or more interfaces from each of the leaf switches to associate to the uplinks. If you want to split a port speed, breakout the interface first before associating the interface to the uplinks. See Breakout Ports to configure breakout the ports from SFS GUI.
   - Select the LAG mode based on the configuration setup in the external network, and click **Next**. To form a LAG on the leaf switches, select an interface or interfaces that are of the same speed.

     ⓘ **NOTE:** Ensure that the corresponding ports on the external switches are configured with the same LAG mode.
5. Associate the networks with the selected interfaces:
   - If the network is already configured, select the untagged networks from the list.
   - If the network is not available, create a network using the **Create Network**—Enter the name, description, and VLAN ID. Networks that are created using this wizard are created on the fabric as general purpose networks. For more information about different types of network in SFS, see Networks.
   - Select **Yes** or **No** appropriately to integrate the networks that are created automatically in the fabric through vCenter on this uplink.

     When you select Yes, the uplink is created with the type `Default` and the networks from the vCenter are automatically appended to the L2 uplink during vCenter integration. For more information, see *OpenManage Network Integration for SmartFabric Services User Guide*.
6. Click **Finish**.

   The system displays user uplink creation success message.

Configure L3 VLAN uplink

Use the following procedure to create a L3 VLAN uplink:

1. Click **Uplink** > **Create**.
2. Enter the uplink details:
   - Select the uplink type as **Layer 3**.
   - Select the L3 type as **L3 VLAN**.
   - Enter the name and description (optional) for the L3 VLAN uplink.
3. Click **Next**.

   ⓘ **NOTE:** You can create L3 uplinks on both leaf and spine switches.
4. Associate the interfaces of the spine or leaf switches with the L3 uplink:
   - Select the switch group.
     - **Spine**—Select a spine switch and an interface or multiple interfaces of the spine switch to be associated with the uplink.
     - **Leaf**—Select a leaf switch from the rack, and an interface or multiple interfaces of the leaf switch to be associated with the uplink.
   - Select the static or dynamic LAG based on the configuration setup in the external network.
5. Click **Next**.

6. Associate the networks with the selected interfaces:
   - Select if the network is a tagged or an untagged network.
   - Enter the name, description (optional), VLAN ID, IP addresses, and prefix length.
7. Define a routing policy to associate with the uplink based on the external network connectivity setup.
   - **Static Route**—A route policy template that contains a policy name, description (optional), the network address, prefix length, and next hop IP address.
   - **eBGP**—A routing policy template that contains a policy name, description (optional), the BGP peer IP address, the peer ASN, and BFD configuration.

     NOTE: The network configurations reflect in the switch only after associating the network with an uplink or server profile.
8. Click **Finish**.

   The system displays uplink creation success message.

Configure L3 Routed uplink

Use the following procedure to create a L3 routed uplink:

1. Click **Uplink** > **Create**.
2. Enter the uplink details:
   - Select the uplink port type as **Layer 3**.
   - Select the L3 type as **L3 Routed Interface**.
   - Enter the name and description (optional) for the L3 routed uplink.
3. Click **Next**.
4. Associate an interface from the spine or leaf switch with the L3 Routed uplink.
   - Select the switch group.
     - **Spine**—Select a spine switch and an interface from the spine to associate with the uplink.
     - **Leaf**—Select the rack, switch, and an interface from the leaf switch to associate with the uplink.

       NOTE: You can select only one interface from the spine or leaf switch for L3 Routed uplink.
5. Click **Next**.
6. Associate the networks with the selected interfaces by providing a name, description (optional), interface IP address, and prefix length.

   NOTE: You cannot associate a L3 Routed network with more than one uplink or server profile.
7. Define a routing policy to associate with the uplink based on the external network connectivity setup.
   - **Static Route**—A route policy template that contains a network prefix and the next hop IP address.
   - **eBGP**—A routing policy template that contains BGP peer IP address, peer AS number, and BFD configuration.
   - Select the **Fabric Interconnect** checkbox, if you are creating the uplink for multisite fabric interconnect purpose. See Multisite fabric interconnect.
8. Click **Finish**.

   The system displays uplink creation success message.

Edit networks and ports for an uplink

Select the uplink from the displayed list to view the details of the uplink on the right. Edit the network and port configuration for an uplink using the following instructions:

Edit networks

1. Select the uplink from the list, and click **Edit Networks**.
2. Edit the untagged network that is associated with the uplink.
3. Click **Update**.

   The system displays the uplink network edit success message.

Edit ports

1. Select the fabric uplink from the list and click **Edit Ports**.
2. Edit the ports that are associated with uplink interfaces.
3. Click **Update**.

The system displays the uplink interface edit success message.

Delete an uplink

You can delete a user-created uplink:

1. Select the uplink from the list and click **Delete**.
2. Click **Delete** to confirm.

   ⓘ **NOTE:** When you delete an uplink, the network and route profile that are associated with the uplink are not deleted.

   The system displays the uplink interface edit success message.

Multisite fabric interconnect

From 10.5.3.0 release, SFS allows you to configure fabric interconnect to connect fabrics across different sites. Using this feature, you can configure the uplink, L3 routed network, IPv4 BGP peer policy, and BGP EVPN policy to interconnect the SFS fabric connected through a direct link or an external network. You can configure multisite fabric interconnect using **Configure Multi-Site Fabric Interconnect** wizard or you can individually create uplink, L3 routed network, and IPv4 BGP peer and BGP EVPN policies for the fabrics in different sites depending on the requirement.

SFS provides two options (direct and external) to configure multisite fabric interconnect depending on the interconnect link between the SFS fabrics. You can configure fabric interconnect using **Multi-Site Fabric Interconnect** option for the supported fabric interconnect links:

1. Direct—Following image is an example of direct interconnect between the multisite SFS fabrics.
   - Direct link between the leaf switches across two sites:

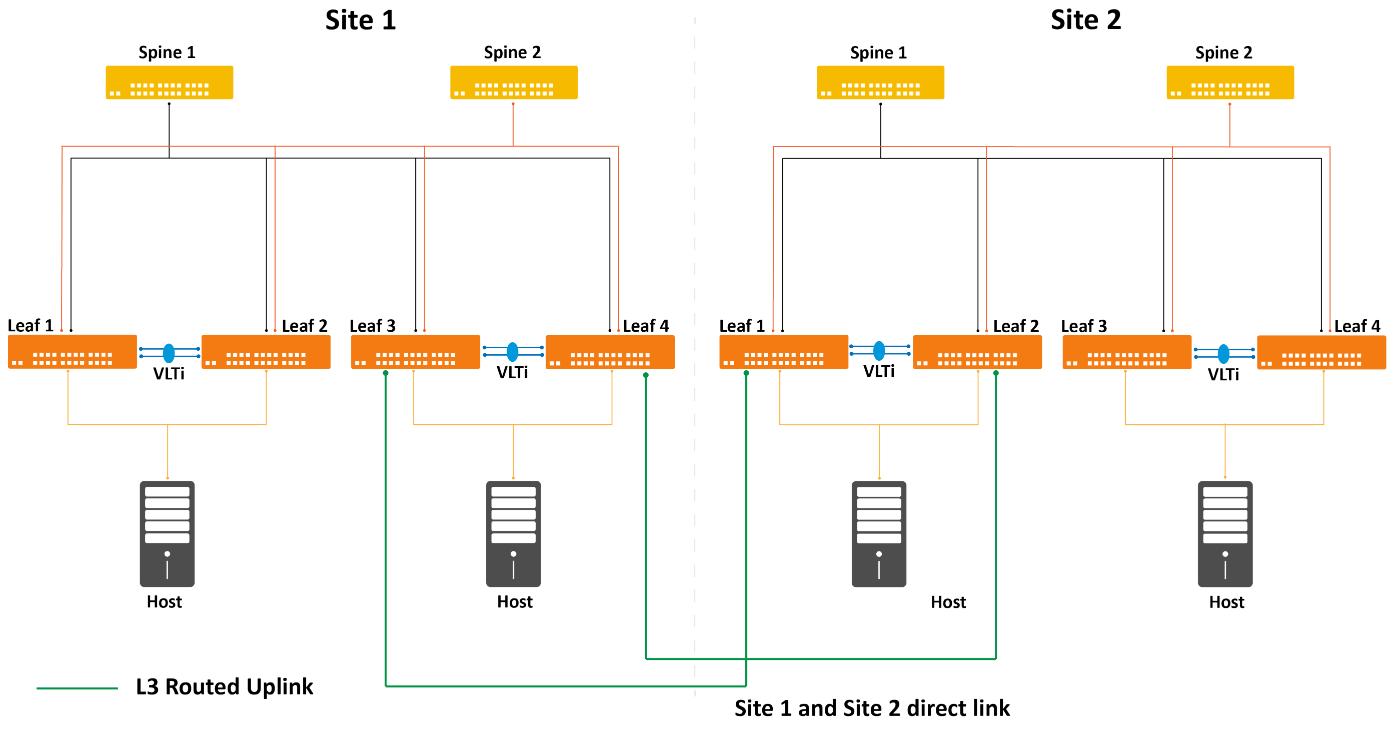

- Direct link between the spine switches across two sites:

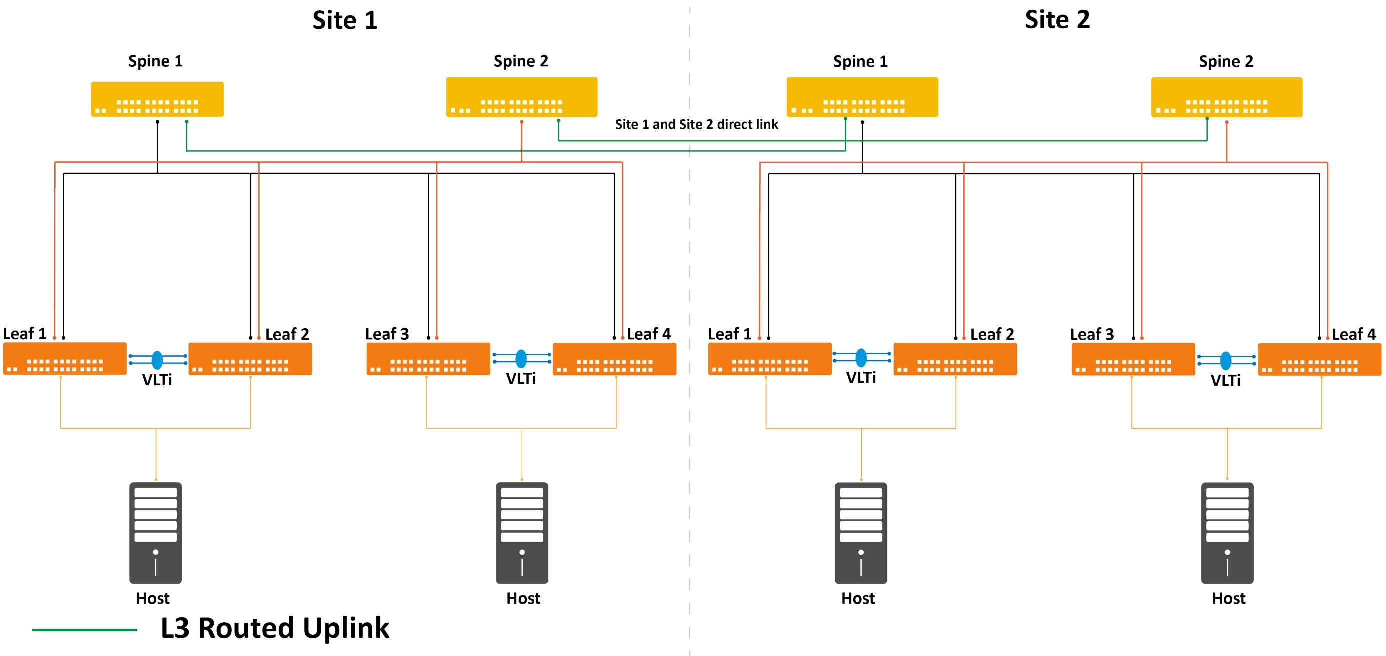

2. External link—Following image is an example of external network interconnect between the multisite SFS fabrics.

- External link between the leaf switches across two sites:

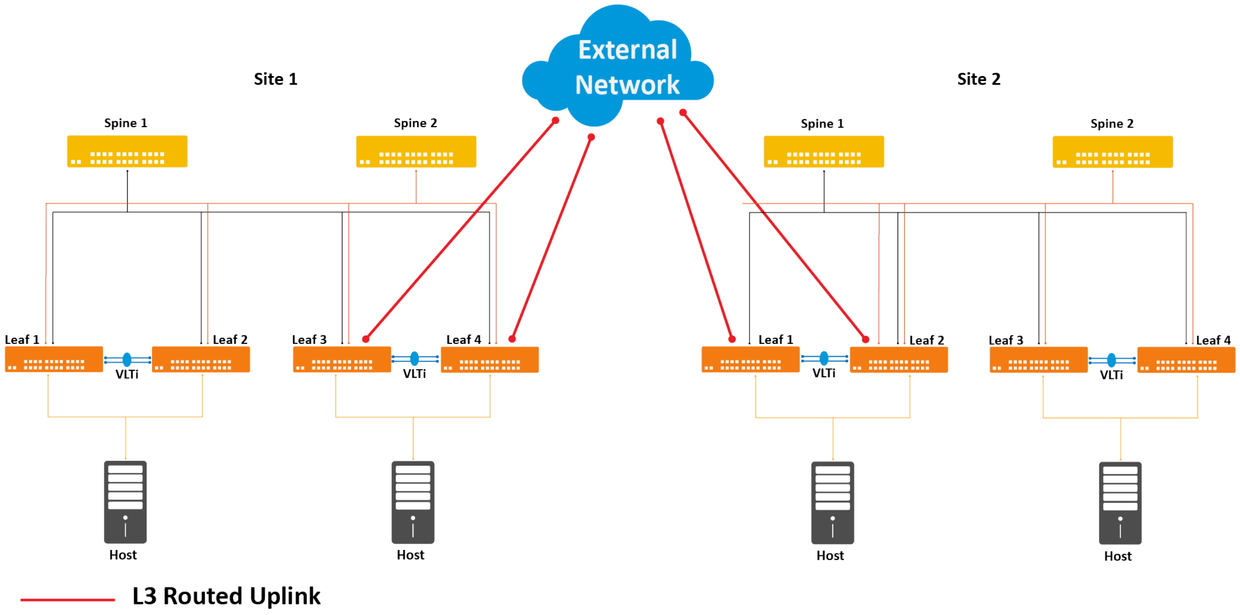

- External link between the spine switches across two sites:

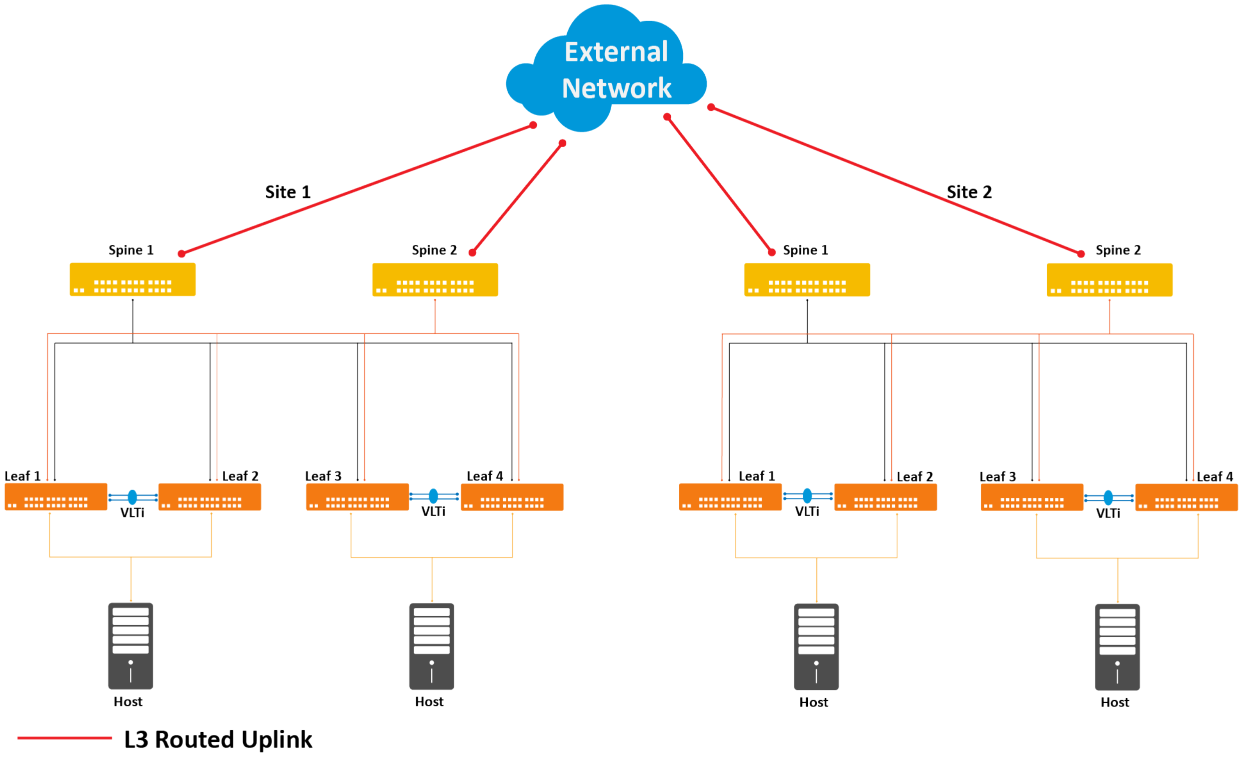

You must configure fabric interconnect configurations individually using Workflow 2:

- When you want to configure the multisite fabric interconnect for leaf and spine switches different from the supported fabric interconnect links.
- When you want to customize the hop-count when creating fabric interconnect.

Prerequisites to setup multisite fabric interconnect

Ensure that the following requirements are met before configuring multisite fabric interconnect:

- Verify the physical connections on the fabric interconnect. SFS supports the fabric interconnects links (direct or external) between leaf switches and spine switches.
- Ensure that the domain IDs of the fabrics are different.
- Ensure that the ASNs of the switches are different.
- Launch the SFS GUI of both the fabrics (Site 1 and Site 2) to collect the following information that you have to enter during the fabric interconnect setup:
    - Domain IDs of the Site 1 and Site 2 fabrics—From SFS GUI, click **Global Settings** > **Fabric Settings** to view the domain ID.
    - BGP ASN numbers of the switches deployed in the Site 1 and Site 2 fabrics—From SFS GUI, click **Global Settings** > **Fabric Settings** to view the ASN details.
    - Loopback address of the switches deployed in the Site 1 and Site 2 fabrics—From SFS GUI, click **Global Settings** > **Fabric Interconnect Details** to view the loopback address of each switch in the fabric.

Procedure to configure multisite fabric interconnect

Following instructions provides the high-level procedure to configure fabric interconnect:

1. Provide the domain ID for the SFS fabrics in Site 1 and Site 2. You can do any of the following:
    - Enable SFS L3 personality on the leaf and spine switches with the domain ID. Domain ID is an optional parameter. If you do not provide the domain ID when enabling SFS, it is set to default domain ID 100. For more information, see smartfabric l3fabric enable. Ensure that the VxRail Manager Integration option is set to **Disable**.
    - Edit the domain ID for the SFS fabric if necessary. Domain ID ranges from 100 to 107.

      NOTE: Editing the domain ID reloads all the switches in the fabric.

      If you edit the domain ID, the following values are autogenerated after reload based on SFS recommendations:

    - Leaf and spine ASN
    - Fabric BGP subnet and mask
    - Fabric VTEP subnet and mask
    - VxRail Manager Integration option is set to **Disable**
2. Configure uplink, L3 routed network, IPv4 BGP peer policy, and BGP EVPN policy in the Site 1 fabric using **Multi-Site Fabric Interconnect** option or individually.
3. Configure uplink, L3 routed network, IPv4 BGP peer policy, and BGP EVPN policy in the Site 2 fabric using **Multi-Site Fabric Interconnect** option or individually.

Workflow 1—Configure multisite fabric interconnect

You can configure the uplink, L3 routed network, IPv4 BGP peer route policy, and BGP EVPN policy to interconnect the SFS fabric connected through a direct link or an external network. Use the following instructions to configure the multisite fabric interconnect:

Configure multisite fabric interconnect—Site 1 fabric

When configuring the peer details for Site 1, provide the switch details of Site 2 fabric as the peer configuration.

1. Click **Uplink** > **Multi-Site Fabric Interconnect**.
2. Click **Configure Multi-Site Fabric Interconnect**.
3. Select the switch group (Leaf or Spine). The template changes based on the selected switch group.
4. If you select the switch group as Leaf, select the rack from the list. The local fabric domain ID, local hostname, local BGP ASN, local loopback address are autopopulated.
5. Select the Fabric Interconnect Link (Direct or External). The template changes based on the selected link type.
6. Enter the domain ID of the peer fabric (Site 2). The system autopopulates the peer BGP ASN number based on the peer domain ID.

   ⓘ **NOTE:** The domain ID of the peer fabric (Site 2) must be different.
7. Enter the following information:
   - Select the uplink interface of the switches that are connected to the respective peer-site switch interfaces.
   - Configure a Layer 3 routed network for the uplink.
     - **Local IP Address**—Enter the IP address for the selected interface.
     - **Local Prefix Length**—Enter the prefix details.
   - Configure the BGP peer policy for Site 1 fabric.

     If the link is **Direct**:

     - **Peer BGP ASN**—Verify the peer BGP ASN number . The system autopopulates the peer BGP ASN number based on the peer domain ID. If the BGP ASN number is different from the autopopulated number, change the ASN number accordingly.
     - **Peer Interface IP**—Enter the peer interface IP address.
     - **BFD**—Select the BFD configuration. By default, BFD is disabled.

     If the link is **External:**
     - **Peer BGP ASN**—Verify the peer BGP ASN number . The system autopopulates the peer BGP ASN number based on the peer domain ID. If the BGP ASN number is different from the autopopulated number, change the ASN number accordingly.
     - **Peer Interface IP**—Enter the peer interface IP address.
     - **BFD**—Select the BFD configuration. By default, BFD is disabled.
     - **External ASN**—Enter the ASN number of the external device.
     - **External Interface IP**—Enter the interface IP address of the external device.
   - Configure the BGP EVPN policy for Site 1 fabric.
     - **Peer Loopback**—Enter the peer loopback address.
     - You can add up to two BGP EVPN policies for each switch. Click **+** to add a second loopback address.

     ⓘ **NOTE:** If the switch group is Leaf, you must enter all the relevant details for the listed leaf switches.

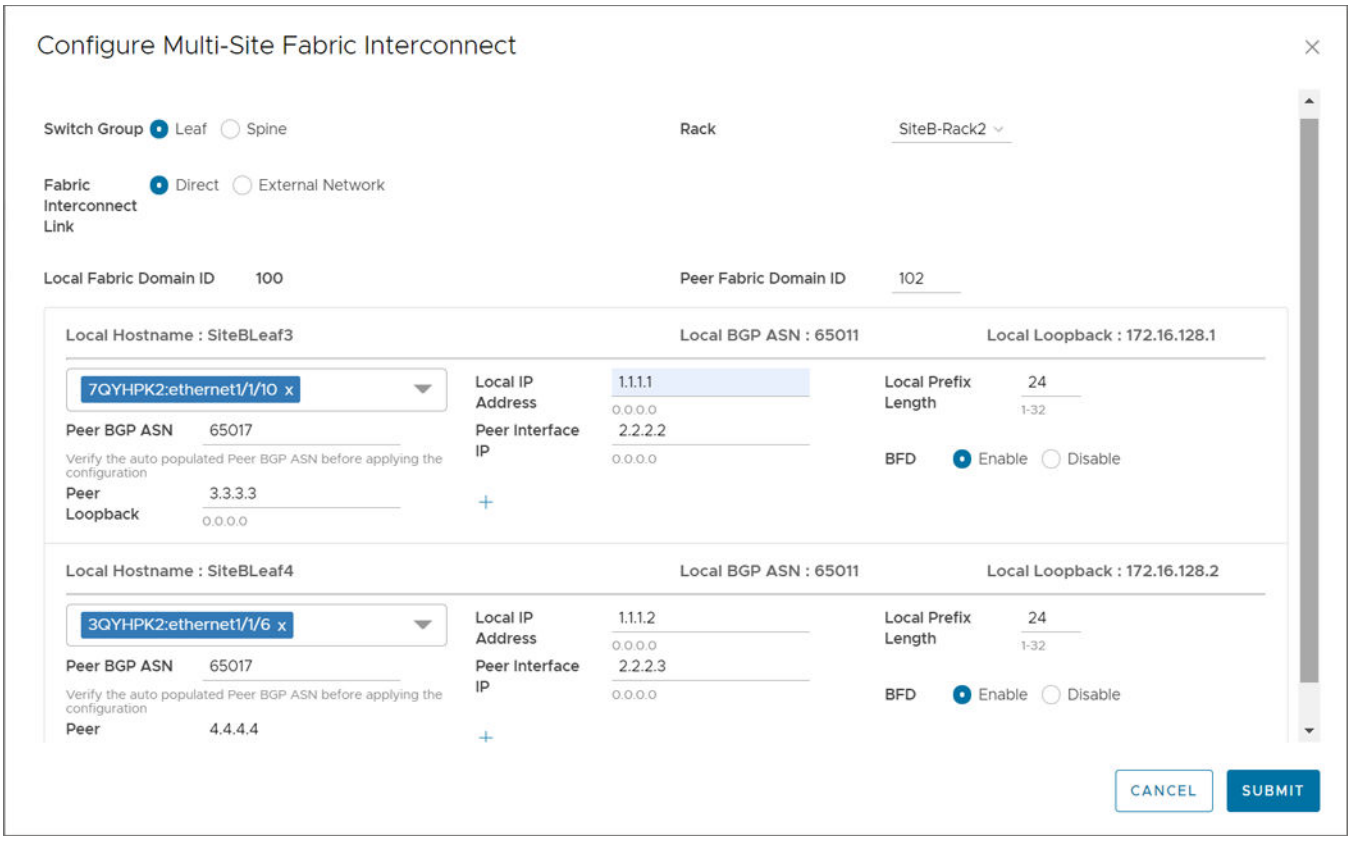

Configure Multi-Site Fabric Interconnect

Switch Group ● Leaf ○ Spine

Rack SiteB-Rack2

Fabric Interconnect Link ● Direct ○ External Network

Local Fabric Domain ID 100

Peer Fabric Domain ID 102

Local Hostname : SiteBLeaf3 Local BGP ASN : 65011 Local Loopback : 172.16.128.1

7QYHPK2:ethernet1/1/10

Local IP Address 1.1.1.1

Local Prefix Length 24

Peer BGP ASN 65017

Verify the auto populated Peer BGP ASN before applying the configuration

Peer Interface IP 2.2.2.2

BFD ● Enable ○ Disable

Peer Loopback 3.3.3.3

Local Hostname : SiteBLeaf4 Local BGP ASN : 65011 Local Loopback : 172.16.128.2

3QYHPK2:ethernet1/1/6

Local IP Address 1.1.1.2

Local Prefix Length 24

Peer BGP ASN 65017

Verify the auto populated Peer BGP ASN before applying the configuration

Peer Interface IP 2.2.2.3

BFD ● Enable ○ Disable

Peer 4.4.4.4

8. Click **Submit**.

After submitting the configuration, SFS configures the following:

- Enables fabric interconnect
- Disables VxRail Manager Integration
- Creates a L3 routed network uplink
- Creates a BGP peer route policy
- If the link is direct, it creates a BGP EVPN policy with the hop count set to 3
- If the link is external, it creates a BGP EVPN policy with the hop count set to 4

You cannot change the hop count in this workflow (Workflow 1). If you want to customize the hop count, use Workflow 2.

The uplink, network, the eBGP peer route, and BGP EVPN policies are created, and their names are appended with the prefix **Fabric_Interconnect**. To verify the configurations, you can go to the respective configuration menus to view the list of uplinks, L3 routed networks, BGP peer route policy, and BGP EVPN policy that are created using this wizard.

Configure multisite fabric interconnect—Site 2 fabric

To configure multisite fabric interconnect on Site 2 fabric, follow the same steps to configure the L3 routed network, IPv4 BGP peer route policy, and BGP EVPN policy. While configuring, provide the switch details of Site 1 fabric as the peer configuration.

Workflow 2—Configure uplinks, L3 routed networks, and BGP peer and EVPN policies

Alternative to the Workflow 1—**Configure multisite fabric interconnect**, you can also configure multisite fabric interconnect by creating the configurations individually. Use this workflow to configure the multisite fabric interconnect for leaf and spine switches that are different from the supported fabric interconnect links:

1. Disable VxRail Manager Integration option, see Edit fabric settings.
2. Create L3 routed interface uplink with eBGP peer route policy configuration, see Uplink. Ensure that you select the **Fabric Interconnect** checkbox when creating the L3 routed uplink.

   NOTE: You can also create eBGP peer route policy individually and associate with the uplink. Ensure that you select the **Fabric interconnect** checkbox when creating the eBGP peer configuration.
3. Create BGP EVPN configuration, see Configure BGP EVPN. Ensure that you provide the appropriate **Hop Count**.

Configure networks and routing configuration

You can configure networks and routing configuration.

Configure networks

You can configure the following types of networks:

- General purpose network
- L3 routed interfaces (for L3 profiles only)
- Multi rack L3 VLAN
- VLAN networks (for L2 and L3 profiles)
- VXLAN networks (for L2 and L3 profiles)

For detailed information about these network types, see Networks.

Configure general purpose networks

When you create a general purpose network, SFS creates a VLAN network along with the VXLAN virtual network.

In general purpose network, VXLAN network identifier (VNI) and VLAN ID are same and you can associate one VLAN with the VNI across the fabric. If you delete a VLAN network, it automatically deletes the associated VXLAN network.

For example, if you create a general purpose network with VLAN ID 50, SFS creates a VLAN 50 and associated VXLAN network with VNI 50 in the SmartFabric. When you delete the VLAN network, both VLAN 50 and VXLAN VNI 50 are deleted.

Create general purpose network

To create a general purpose network:

1. Click **Network** > **Networks** > **General Purpose Networks**. The page displays the list of the general purpose networks that are already configured in the SmartFabric.
2. Click **Create** to create a Layer 2 general purpose network.
3. Enter the following details:
   - Network ID
   - Network name
   - VLAN ID. The number ranges from 1—3999 (excluding 3939)
   - (Optional) Description
4. Click **Create**. The system displays virtual network creation successful message.

View general purpose network

The general purpose networks display the VLAN details of the specific network including network ID, network name, VLAN ID, QoS priority, and network type. `Portgroups` that are created on the vCenter are displayed under **General Purpose Networks**.

Edit general purpose network

You can edit the configuration of the Layer 2 general purpose network and change it to Layer 3 general purpose network.

1. Select a network from the list and click **Edit**.
2. Modify the Network Type.
3. Enter the Prefix, Gateway IP Address, IP address.
4. Click **Edit**. The system displays virtual network edits success message.

Delete general purpose network

When you delete a general purpose network, both the VLAN and the VXLAN networks are deleted. To remove a general purpose network configuration:

1. Select the general purpose network, and click **Delete**. The system displays the list of the server interface profiles associated with the network.
2. Click **Delete** to confirm. The system displays network deletion success message.

Configure L3 routed interfaces

This information explains how to create, edit, and delete Layer 3 routed interfaces.

Create L3 routed interface

Use the following procedure to create an L3 routed interface:

1. Click **Network**.
2. Select **Networks** > **Layer 3 Routed Interfaces**.
3. Click **Create**.
4. Enter the following details:
   - Network ID
   - Network name
   - (Optional) Description
   - IP Address
   - Prefix length
5. Click **Create**.

   The system displays network creation success message.

Select any network from the list to view the detailed information of the routed interface network.

Edit network

1. Select the **Network ID** from the list and click **Edit**.
2. Edit the configuration as required.
3. Click **Edit**.

   The system displays edit network success message.

Delete network

1. Select the network ID to remove and click **Delete**. The system displays the list of uplinks associated with the network.
2. Click **Delete** to confirm.

   The system displays network deletion success message.

Configure multirack L3 VLAN

Starting from 10.5.3.0 release, from SFS GUI, you can configure L3 VLAN network for the racks to which the servers are connected. Using this feature, you can create a L3 VLAN network for each VLT pair (rack) with a different subnet. This network is used for NSX-T overlay to create VTEP networks. Create, edit, and delete multirack L3 VLAN networks from SFS GUI. You can specify IP address for each switch in a rack when creating a multirack L3 VLAN.

Create multirack L3 VLAN

1. Click **Network**.
2. Click **Networks** > **Multi-Rack L3 Networks**.
3. Click **Create**.
4. Enter the following details:
   - Network ID
   - Network name
   - VLAN ID
   - (Optional) Description
   - Select the **Specific IP Addresses** checkbox to specify the IP address for each switch.
   - IP addresses
   - Prefix length
   - Gateway IP address
   - Helper IP address
5. Click **Create**.

   The system displays VLAN network creation success message.

Select any network from the list to view the detailed information.

Edit multirack L3 VLAN configuration

1. Select a network ID you want to edit and click **Edit**.
2. Edit the details and configuration as necessary.
3. Click **Edit**.

   The system displays edit network success message.

Delete multirack L3 VLAN configuration

1. Select the VLAN network to remove and click **Delete**.
2. Click **Delete** to confirm.

   The system displays network deletion success message.

Configure VLAN networks

Create, edit, and delete L2 or L3 VLAN networks for SmartFabric.

Create L2 VLAN

VLAN networks for L2 profile:

1. Click **Network**.
2. Select **Networks** > **VLAN Networks**.
3. Click **Create**.
4. Select the network type as **Layer 2 VLAN Network**.
5. Enter the following details:
   - Network ID
   - Network name
   - VLAN ID
   - (Optional) Description
6. Click **Create**. The system displays VLAN network creation success message.

VLAN networks for L3 profile:

1. Click **Network**.
2. Select **Networks** > **VLAN Networks**.
3. Click **Create**.
4. Select the network type as **Layer 3 VLAN Network**.
5. Enter the following details:
   - Network ID
   - Network name
   - VLAN ID
   - (Optional) Description
   - IP addresses
   - Prefix length
   - Gateway IP address
   - Helper IP addresses
6. Click **Create**.

   The system displays VLAN network creation success message.

Edit network

You can edit any changes to the existing configuration. Using this option, you can also change the network types such as from L2 to L3 or L3 to L2.

1. Select a network ID from the list and click **Edit**.
2. Modify the configuration as necessary.

3. Click **Edit**.

   The system displays edit network success message.

Delete network

1. Select the VLAN network to remove and click **Delete**.
2. Click **Delete** to confirm.

   The system displays network deletion success message.

Configure VxLAN network

Create, edit, and delete L2 and L3 profile VXLAN network configurations through SFS GUI. The purpose of VXLAN network is to associate multiple L2 or L3 VLAN networks to a single VXLAN network. Whereas a general purpose network does not have the flexibility to extend the VXLAN network.

Create VxLAN network

Virtual network for L2 profile:

1. Click **Network**.
2. Select **Networks** > **VxLAN Networks**. The page displays the list of the VXLAN networks that are configured in the SmartFabric instance.
3. Click **Create**.
4. Select **Layer 2** as the virtual network type.
5. Enter the following details:
   - Virtual network name
   - VxLAN VNI
   - VLT VLAN ID
   - (Optional) Description
6. Click **Create**.

   The system displays virtual network creation successful message.

Virtual network for L3 profile:

From 10.5.3.0 release, you can configure the IPv6 virtual networks.

1. Select **Networks** > **VxLAN Networks**. The page displays the list of the VXLAN networks that are configured in the SmartFabric instance.
2. Click **Create**.
3. Select **Layer 3** as the **Virtual Network Type**.
4. Enter the following:
   - Virtual network name
   - VxLAN VNI
   - VLT VLAN ID
   - Select IPv4 checkbox to provide the IPv4 address, prefix, gateway IP address, and helper addresses.
   - Select IPv6 checkbox to provide the IPv6 address, prefix, gateway IP address, and helper addresses.
5. Click **Create**.

View VxLAN network details

The VxLAN networks display a list of mapped VLANs. Select a VxLAN network to view details pertaining to that specific network including network ID, VLAN ID, and network name.

Associate multiple VLANs to a VxLAN network

Using the following steps, you can map multiple VLANs to a single VxLAN network:

1. Select a VxLAN network.
2. Click **Create** option available after the VxLAN details.
3. Enter the required details for the VLAN configuration.
4. Click **Create**.

Edit VxLAN network

You can edit the configuration of VXLAN network configuration. You can also change the network type such as from L2 to L3 or L3 to L2.

1. Select a virtual network from the list and click **Edit**.
2. Modify the network type.
3. Enter the relevant details.
4. Click **Edit**.

   The system displays virtual network edits success message.

Delete VxLAN network

To delete a VXLAN network, first delete the mapped VLAN or VLANs if associated, and delete the virtual network.

1. Select the Virtual Network Name, select the Network to remove, then click **Delete**.
2. Click **Delete** to confirm.

   The system displays network deletion success message.

Configure Routes

You can configure static routes, eBGP peer routes, and BGP EVPN configuration for a network.

When you create a route policy, the system configures the route policy name as the policy ID. In SFS, the route policy ID must be unique and hence it does not allow you to create more than one policy with the same name. You can create multiple route policies with the same configuration by providing different route policy name.

You can create and delete any routing policy configuration, but you cannot edit any routing policy. To modify an existing routing policy, you must delete the policy and create a policy with the relevant configuration.

Configure static routes

Configure static routes and associate the route to the switch.

Create static route

1. Click **Network** > **Routing Configuration**.
2. Select **Static Routes**.
3. Click **Create** to add a new static route.
4. Enter the following details:
   - Policy name
   - Network address
   - Prefix length
   - Next hop IP address
   - (Optional) Description
5. Click **Create**.

   The system displays static route creation is successful.

Add route to switch

You can add the routing policy to the switch:

1. Select a static route that must be added to the switch.
2. Click **Add Route to Switch**.
3. Select the switch to map to this route.
4. Click **Add**.

   The system displays the route added success message.

The static route details display a list of mapped routes. Select a static route to view details pertaining to that specific route including the switch ID.

Delete route from switch

1. Select the route to delete, and click **Delete Route**.
2. Click **Delete** to confirm the removal of the route from the switch.

   The system displays route policy deletion success message.

Delete static route

1. Select the static route to delete and click **Delete**.
2. Click **Delete** to confirm.

   The system displays static route deletion is successful.

Configure eBGP peer route

You can configure eBGP peer routes for a network.

Create eBGP route

1. Click **Network** > **Routing Configuration**.
2. Select **eBGP Peer Configuration**.
3. Click **Create**.
4. Enter the following details:
   - Policy name
   - Peer interface IP address
   - Peer ASN
   - (Optional) Description
   - BFD configuration
   - Select the **Fabric Interconnect** checkbox, if you are creating the uplink for multisite fabric interconnect purpose. See Multisite fabric interconnect.
5. Click **Create**. The system displays eBGP peer route creation is successful.

The eBGP peer details display a list of mapped routes. Select an eBGP route to view details pertaining to that specific route including the switch ID.

Delete eBGP route

1. Select the eBGP route policy to delete and click **Delete**.
2. Click **Delete** to confirm. The system displays route policy deletion success message.

Add eBGP route to switch

1. Select an eBGP route policy from the list.
2. Click **Add Route to Switch**.
3. Select the switch to which you want to add the route.
4. Click **Add**. The system displays the route to switch addition success message.

Delete eBGP route from switch

1. Select an eBGP route, then click **Delete Route**.
2. Click **Delete** to remove the route from the switch. The system displays route deletion success message.

Configure BGP EVPN

Follow the instructions to configure a BGP EVPN routing policy:

Create BGP EVPN configuration

1. Click **Network** > **Routing Configuration**.
2. Click **BGP EVPN Configuration**.
3. Click **Create**.
4. Enter the following details:

- BGP EVPN policy name
- Peer loopback address
- Peer ASN
- Hop count
- (Optional) Description

5. Click **Create**. The system displays eBGP peer route creation is successful.

Delete BGP EVPN configuration

1. Select a BGP EVPN configuration and click **Delete**.
2. Click **Delete** to remove the route from the switch. The system displays route deletion success message.

Add BGP EVPN route to switch

1. Select a BGP EVPN route policy from the list.
2. Click **Add Route to Switch**.
3. Select the switch to which you want to add the route.
4. Click **Add**. The system displays the route to switch addition success message.

Delete eBGP route from switch

1. Select a eBGP route and click **Delete Route**.
2. Click **Delete** to remove the route from the switch. The system displays route deletion success message.

Configure global settings for SmartFabric

Starting from 10.5.3.0 release, you can configure SmartFabric switch services settings using SFS GUI.

You can configure the following services on the SmartFabric switches:

- NTP
- DNS
- Syslog
- SNMP

ⓘ **NOTE:** This feature is supported from SmartFabric OS10.5.2.2 and later versions, and applicable for SFS L3 leaf and spine personality.

Configure NTP server

To configure an NTP server:

1. Click **Global Settings** > **NTP**.
2. Click **Add** to configure an NTP server.
3. Enter the IP address or hostname of the NTP server and click **Add**.

   The system displays the configuration success message.

To delete an NTP server, select an entry from the list and click **Delete**.

Configure DNS server

To configure one or more DNS servers:

1. Click **Global Settings** > **DNS**.
2. Click **Add** to configure one or more DNS servers.
3. Enter the IP address of the DNS server to configure a single DNS server setting. You can use the **+** button to add more DNS servers.
4. Click **Add**.

   The system displays the configuration success message.

To delete the configured servers, select the server from the list and click **Delete All**. This action deletes all the configured DNS servers that are available in the system.

Configure SNMP server

To configure or edit an SNMP server:

1. Click **Global Settings** > **SNMP**.
2. Click **Add** to configure an SNMP server.
3. Enter the following details:
    - Server IP address
    - Community
    - SNMP version
    - Provide the **Security** details for SNMP v3.
4. Click **Add**.

    The system displays the configuration success message.

To delete the configured servers, select a server from the list and click **Delete**.

Configure syslog server

To configure and edit a syslog server:

1. Click **Global Settings** > **Syslog**.
2. Click **Add** to configure syslog server.
3. Enter the IP address of the syslog server and log level.
4. Click **Add**.

    The system displays the configuration success message.

Edit syslog server

You can edit the log level for the syslog server.

1. Select the server from the list and click **Edit**.
2. Edit the log level of the server and click **Update**.

To delete the configured servers, select the server from the list and click **Delete**.

Edit default fabric settings

Use the following procedure to edit the default fabric settings:

1. Click **Global Settings** > **Fabric Settings**.
2. Click **Edit**.
3. Change the following values of the default settings as required:
    - Domain ID
    - Leaf and spine ASN
    - Fabric BGP subnet and mask
    - Fabric VTEP subnet and mask
    - Client control VLAN
    - Client management VLAN
    - STP mode between MST and Rapid PVST
    - VxRail Manager Integration
4. Click **Submit**. The changed settings are applied only after a reboot. The system prompts for confirmation to continue. After you click **Ok**, all the switches in the network fabric reload to apply the changes.

When you edit the domain ID:

1. The fabric settings, multisite fabric interconnect configuration, and fabric name are changed after reload.
2. The interface breakout, networks, server profiles, and uplink configurations and its association are retained after reload.
3. You must reconfigure BGP policy details for the external peer connectivity configurations such as multisite fabric interconnect links, external router, NSX-T based on the changed ASN.

To manually onboard the VxRail node onto the fabric:

1. Click **Global Settings** > **Fabric Settings**.

2. Click **Edit**.
3. Select **Disable** for VxRail Manager Integration.

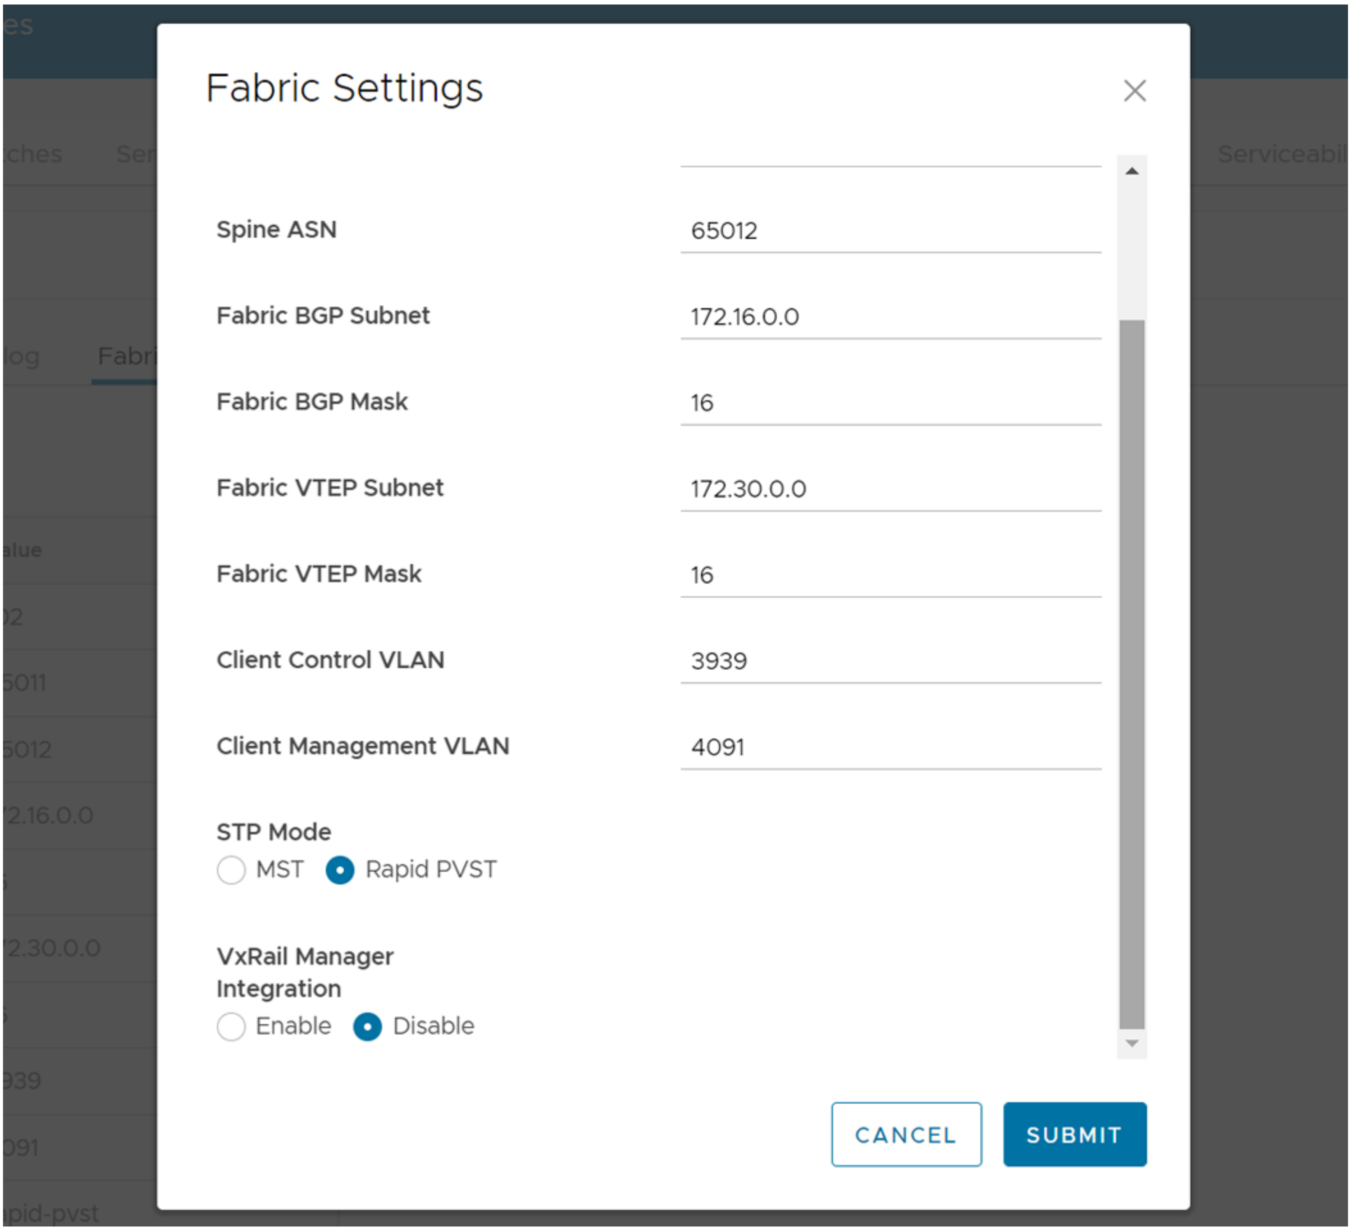

Fabric Settings

Spine ASN	65012
Fabric BGP Subnet	172.16.0.0
Fabric BGP Mask	16
Fabric VTEP Subnet	172.30.0.0
Fabric VTEP Mask	16
Client Control VLAN	3939
Client Management VLAN	4091

STP Mode
○ MST ● Rapid PVST

VxRail Manager Integration
○ Enable ● Disable

4. Click **Ok**.

   ⓘ **NOTE:** The VxRail Manager Integration setting change does not reload the switches, and the change is applied to the latest selection immediately.

After disabling the VxRail Manager Integration, you can onboard the VxRail server to the fabric manually. If you want to dynamically onboard the new VxRail nodes, you can enable the VxRail Manager Integration option.

Update default fabric, switch names, and descriptions

SFS assigns unique names for the network fabric, racks, and switches automatically. Use the following instructions to change the names and descriptions of the network fabric, racks, and switches:

1. Click **Global Settings** > **Set Fabric & Switch Name**.
2. Change the name and description of the network fabric, and click **Next**.
3. Change the name and description of the rack or VLT fabric, and click **Next**.
4. Change the name and description of the switches, and click **Finish**.

ⓘ **NOTE:** If you change the switch name in the GUI, the hostname on the switch CLI is updated.

View fabric interconnect details

View the fabric interconnect details:

Click **Global Settings** > **Fabric Interconnect details**. This page displays the hostname of the switches and the respective loopback address of the switches in the fabric. You can use this page to view the loopback address when configuring multisite fabric interconnect.

Restore fabric configuration

Restore the SFS configurations on a fabric to a known good configuration with the backup configuration file stored on your external device. Use the following procedure to restore:

1. Click **Life Cycle Management** > **Restore**.
2. Click **Choose File**.
3. Select the backup configuration json file that is stored externally
4. Click **Ok**.
5. Select the check box to agree.

   ⚠ CAUTION: **This action reboots all the switches in the fabric to apply the new configuration that cause traffic disruption and connectivity. Dell Technologies recommends you to use the restore operation during a downtime window.**
6. Click **OK** to confirm. The switches are restored with the configuration in the JSON file.

View fabric compliance status

From 10.5.3.0, you can view the fabric compliance status from SFS GUI. SFS validates the health of the cluster, topology role, underlay, overlay, network, server appliance discovery, uplink, policy, and VLT. SFS monitors the health in both the switch and the whole fabric levels. SFS retrieves the fabric compliance status for the SFS instance and displays the noncompliance events with details. SFS also provides recommended actions to eliminate the compliance violations or misconfigurations.

This feature is applicable for SFS L3 leaf and spine personality.

To view the fabric compliance errors:

1. Click **Serviceability** > **Fabric Compliance** to view the latest compliance errors. The table lists the latest compliance events with detailed information including switch name, service tag of the switch, status, error code, and the recommended action.
2. Click the information icon to view the recommended action for each compliance error.
3. Click **Refresh** to update the data and display the new compliance errors.
4. Click **Download** to download all the compliance errors. The downloaded `zip` file contains the fabric compliance errors in CSV format.

You can also view the fabric compliance errors in the SmartFabric instance overview dashboard. Select **Summary** > **Overview** to view the overview of fabric compliance status. The fabric compliance errors are grouped under different categories.

6

SFS with VxRail

SFS, used in leaf and spine network, creates a fully integrated solution between the fabric and a hyperconverged domain infrastructure such as VxRail. When integrated with VxRail, SFS automates network setup, simplifying and accelerating the deployment. Switches are automatically configured. When additional VxRail nodes are connected, the fabric identifies them as VxRail nodes and automatically onboards the nodes to the required networks.

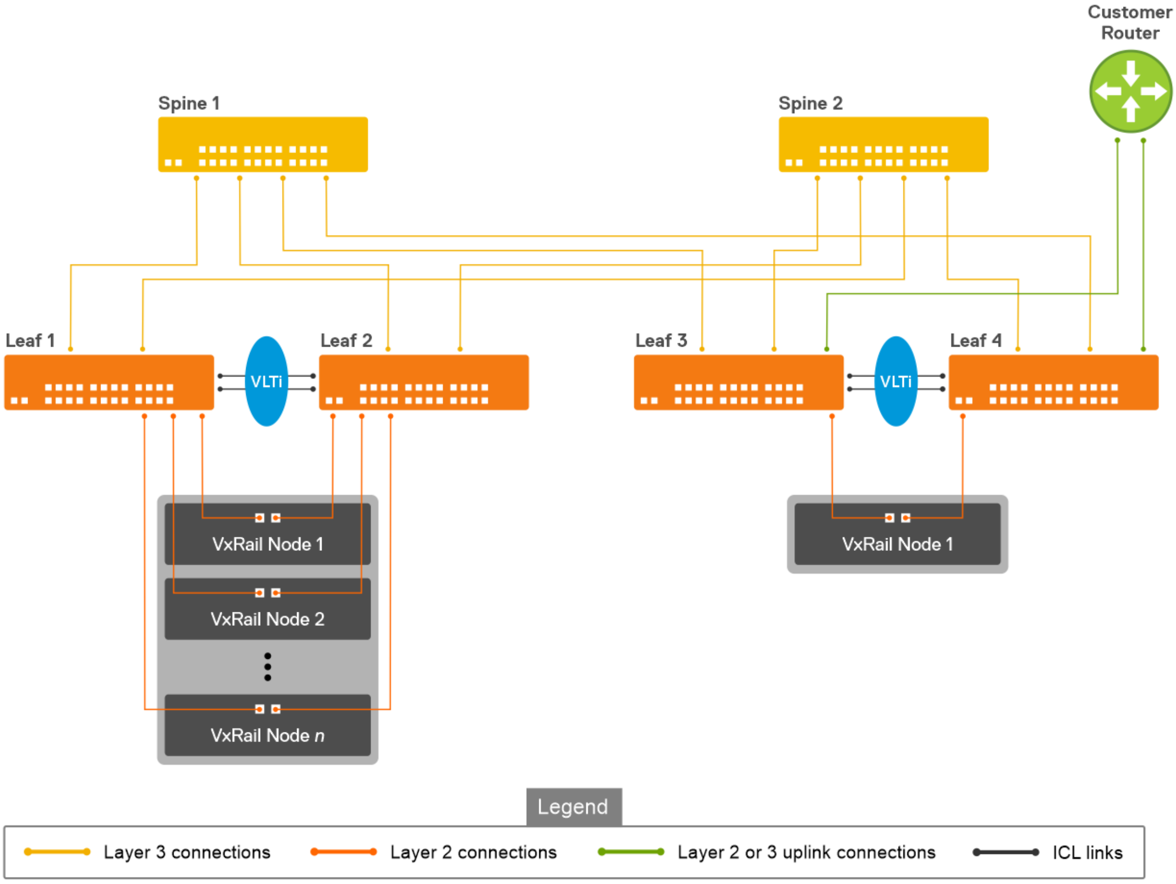

For more information regarding VxRail deployment-related documents, see Dell Technologies VxRail Networking Solutions.

Supported network topologies

See Supported topologies section for information regarding the topologies of SFS with VxRail deployments.

Hardware and software requirements

The requirements to deploy VxRail with SFS are as follows:

Hardware components

- VxRail nodes
- Dell EMC PowerSwitches

Software components

- SmartFabric OS10
- OMNI
- VxRail Manager
- VMware vCenter

For more information regarding detailed deployment requirements, see the Deployment Guides for respective releases.

Supported switches

Following is the Dell EMC PowerSwitch typical roles in SFS with VxRail deployment.

Table 7. Switch roles in SFS

SmartFabric Switches	Switch role (leaf or spine)	VxRail node connectivity options
• S4112F-ON • S4112T-ON • S4128F-ON • S4128T-ON • S4148F-ON • S4148T-ON	Leaf (top of rack) or spine	10GbE
• S5212F-ON • S5224F-ON • S5248F-ON • S5296F-ON		10GbE or 25GbE
S5232F-ON	Spine	Can be used as a leaf switch with ports that are connected to VxRail nodes broken out to 10GbE or 25GbE
• Z9264F-ON • Z9432F-ON	Spine	—

In VxRail deployment, any combination of the leaf and spine switches is possible with the exception that you must deploy leaf switches in pairs. Each leaf switch in the pair must be the same model due to VLT requirements. SFS supports up to 20 switches and eight racks in the fabric.

SFS personalities

In SFS-enabled network, VxRail deployment option include L2 single rack or L3 multirack personalities. The table lists the comparison between L2 and L3 fabric personalities:

Table 8. SFS personalities in VxRail deployment

L2 Single Rack personality	L3 multi rack personality
Single rack network fabric is supported for VxRail clusters. For new SFS deployments, use the L3 leaf and spine fabric personality as the SFS L2 personality is deprecated.	Multi rack data center network fabric is supported that starts with a L3 single rack (L3 fabric profile) and which you can expand to a multi rack solution based on the demand.
Network fabric has two leaf switches in a single rack which you cannot expand.	Network fabric has up to 20 switches in a leaf and spine design that starts with a single rack which you can expand up to eight racks.
All VxRail with SFS deployments from SmartFabric release OS10.4.1.4 to OS10.5.0.5 support configuration with a single pair of leaf switches for VxRail clusters.	All SmartFabric deployments with SmartFabric OS10.5.0.5 or later.
Default uplink and jump host port are created as part of a fabric initialization, which you cannot modify after enabling SFS.	You can create uplinks and jump host port through SFS GUI or OMNI after initial deployment.

Table 8. SFS personalities in VxRail deployment (continued)

L2 Single Rack personality	L3 multi rack personality
Enabled SFS by running a Python script in the OS10 Linux shell.	Enable SFS using CLI, API, or UI.
Existing deployments when upgraded to SmartFabric OS10.5.0.5 continue to run in the L2 fabric profile and L3 fabric capabilities are not available. If you upgrade switches with L2 personality to OS10.5.0.5, SFS operates with the VxRail L2 single rack personality.	Dell Technologies recommends that you enable all new deployments with L3 leaf and spine fabric personality. You cannot upgrade VxRail with SFS deployments to the new L3 leaf and spine fabric personality automatically.

Support matrix

See Networking Solutions Support Matrix regarding support matrices for SFS with VxRail across various releases.

Fabric operations and life cycle management

Dell EMC OpenManage Network Integration (OMNI) enables you to configure and manage SFS-enabled PowerSwitches in different deployments. You can use OMNI application to manage and operate one or more SFS instances either directly using the OMNI UI through a web browser or through a vCenter plug-in. After initial deployments, Dell Technologies recommends that you use OMNI UI to perform all fabric management and life cycle management activities.

For more information about OMNI, see OMNI Documentation.

7

SFS with PowerEdge MX

Dell EMC PowerEdge MX is a unified, high-performance data center infrastructure providing the agility, resiliency, and efficiency to optimize a wide variety of traditional and new emerging data center workloads and applications. In a Dell EMC PowerEdge MX7000 infrastructure, the MX9116n fabric engine and MX5108n Ethernet switch support SFS.

For more information about the SFS and PowerEdge MX including architecture, deployment, configuration, operations and troubleshooting, see the Dell EMC PowerEdge MX Networking Deployment Guide.

8

SFS for Isilon/PowerScale back-end fabric

Dell EMC PowerScale is a scale-out network-attached storage (NAS) platform that supports unstructured data workloads. All PowerScale models are powered by the OneFS operating system. PowerScale uses Dell EMC PowerSwitches to provide the network. SmartFabric OS10 with SFS, for PowerScale back-end fabric automates onboarding and network configuration of PowerScale devices on a L3 leaf and spine fabric. Isilon OneFS interacts with back-end fabric formed by SFS. For more information, see PowerScale Info Hub.

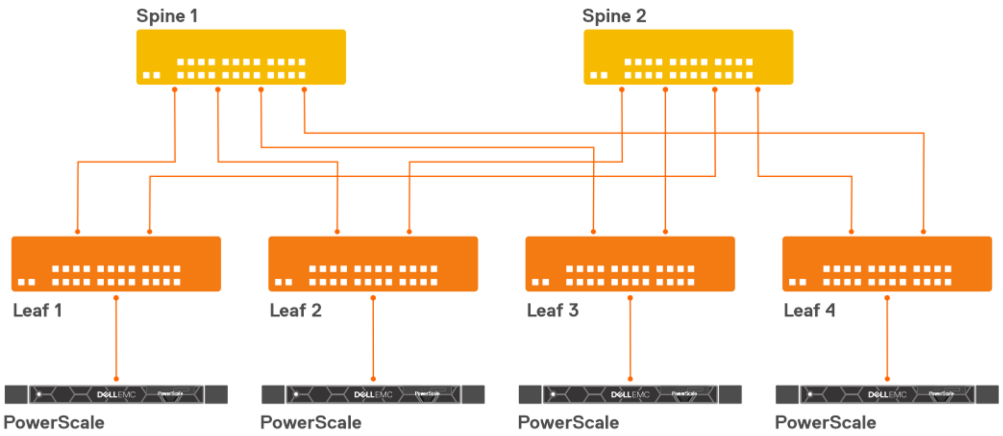

Supported network topologies

See Supported topologies section for information regarding the topologies of SFS. Only L3 personality is supported when deploying SFS with PowerScale nodes.

Hardware and software requirements

The requirements to deploy PowerScale with SFS are as follows:

Hardware components

- PowerScale devices
- Dell EMC PowerSwitches

Software components

- Dell EMC SmartFabric OS10
- Isilon OneFS

Supported switches

The following PowerSwitches are supported in PowerScale deployment:

- S4112F-ON
- S4148F-ON
- S5232F-ON
- Z9264F-ON
- Z9100-ON

PowerScale requirements

Requirements specific to SFS with PowerScale deployment are as follows :

- By default, all PowerSwitches for PowerScale deployment are shipped with factory-loaded OS10.

  > **NOTE:** Dell EMC PowerSwitches must be running SmartFabric OS10.5.0.5 or later software releases that support the PowerScale with SFS deployment.
- In SFS with PowerScale deployment, the leaf nodes are not connected as a VLT pair. On the leaf switch, no ICL configuration is required while enabling SFS. For more information, see Enable SFS using CLI.
- You only have to enable SFS on the switches to setup a fabric. All other fabric operations are managed by Isilon OneFS.
- The switches are connected in a leaf and spine topology with BGP EVPN setup between the leaf switches.
- Leaf and spine switches belong to different autonomous systems (AS).
- PowerScale data traffic is forwarded on an untagged VXLAN network 4091.
- PowerScale appliances are detected through LLDP.
- All ports on which the PowerScale appliance is detected are added as untagged VXLAN access interfaces.
- All load balancing is achieved through BGP ECMP.
- By default, autobreakout feature is enabled on the leaf and spine switches.

Support matrix

See Networking Solutions Support Matrix regarding support matrices for SFS with PowerScale across various releases.

9

SmartFabric commands

You can run `show` commands specific to SFS from the CLI to view fabric configuration information. The command output varies depending on the SFS deployment.

smartfabric l3fabric enable

Enables SFS on the switches and creates a L3 network fabric.

Syntax

```
smartfabric l3fabric enable role {LEAF [vlti ethernet node/slot/port] |
SPINE} [domain domain-ID]
```

Parameters

`role`—Enter the role of the switch in the L3 fabric:

- `LEAF [vlti ethernet node/slot/port]`—Specify the role as `LEAF` for top of rack switches and specify the VLTi ports that interconnect the leaf switches.

  NOTE: Option to specify VLTi ports are not applicable for Isilon/PowerScale deployment.
- `SPINE`—Specify the role as `SPINE` for the switch that connects the leaf switches.
- `[domain domain-ID]`—Specify the domain ID of the switch. The domain ID can range from 100 to 107. By default, domain ID is set to 100.

Default

None

Command Mode

CONFIGURATION

Usage Information

After you enable the L3 fabric and set a role, the system prompts for confirmation to reboot to apply the mode change. If you type `Yes`, the switch reboots in SmartFabric mode and a network fabric is created automatically with default fabric settings. From 10.5.3.0 release, you can enable the L3 fabric along with domain. The leaf and spine ASN and subnet prefixes is autogenerated based on the domain ID.

The `no smartfabric l3fabric` command disables the L3 fabric personality. After you disable the L3 fabric in the switch, the system reboots to change the personality after confirmation.

This command is accessible to users with `sysadmin` and `netadmin` roles.

Example (Spine)

```
OS10(config)# smartfabric l3fabric enable role SPINE domain 101

Reboot to change the personality? [yes/no]: yes
```

Example (Leaf)

```
OS10(config)# smartfabric l3fabric enable role LEAF vlti
ethernet 1/1/4-1/1/5 domain 101

Reboot to change the personality? [yes/no]: yes
```

Example (disable SFS on both leaf and spine)

```
OS10(config)# no smartfabric l3fabric

Reboot to change the personality? [yes/no]: yes
```

Supported Releases

10.5.0.3 or later

smartfabric vlti

Updates the VLTi ports after SFS is enabled.

Syntax

`smartfabric vlti` *`ethernet ports`*

Parameters

`ethernet ports`—Specify the VLTi Ethernet ports which must be updated.

Default

None

Command Mode

CONFIGURATION

Usage Information

Use this command to configure or update the VLTi information after SFS is enabled on the switch. The system reloads with the configured VLTi ports.

This command can be used only if the switch should already be in L3 fabric mode. If not, enable the L3 fabric personality first and run this command.

If you use any of the existing ports for the VLTi, those ports should also be specified as part of the VLTi configuration using the SmartFabric Services commands.

Example

```
OS10(config)#smartfabric vlti ethernet 1/1/31-1/1/32

Warning:The system will be reloaded now, for the personality changes to
take effect
```

Supported Releases

10.5.0.3 or later

show logging smartfabric

Displays important logs that are related to SFS modules.

ⓘ **NOTE:** You must have access to `sysadmin`, `secadmin`, or `netadmin` roles to run this command.

Syntax

`show logging smartfabric`

Parameters

`<1-65535>`—Number of recent messages to be displayed.

Default

Not applicable

Command Mode

EXEC

Usage Information

You can run this command on all switches in SmartFabric mode and are provided with a high-level view of the events happening in the SFS module. When you run this command on an active switch, the system displays complete cluster-related logs. The logs include events that are related to cluster formation and update, configuration changes, and on boarding events along with switch-specific information. When you run this command on a backup switch, the system displays only the switch-specific logs.

Example

```
MX9116N-B2# show logging smartfabric 100 | grep CAGT
2020-03-29 10:09:35.334 MX9116N-B2 [DNV-CAGT] [chassis.get_lead_chassis]
lead chassis not in chassis_data
2020-03-29 10:09:35.335 MX9116N-B2 [DNV-CAGT] [chassis.get_lead_chassis]
not receive mdns from lead MSM yet
2020-03-29 10:09:56.881 MX9116N-B2 [DNV-CAGT]
[app.process_cps_cluster_sync_event] received sync state 2
2020-03-29 10:09:56.881 MX9116N-B2 [DNV-CAGT] [ka.clr_priority]
2020-03-29 10:09:56.885 MX9116N-B2 [DNV-CAGT] [ka.set_priority] new 100
old 0
2020-03-29 10:09:58.014 MX9116N-B2 [DNV-CAGT] [ka.soft_reload_ka]
reloading ka...
```

Supported Releases

10.5.2.3 or later

show smartfabric cluster

Displays the basic cluster information of the switch or IOM.

Syntax `show smartfabric cluster`

Parameters None

Default None

Command Mode EXEC

Usage Information This command is supported in both Full Switch and SmartFabric modes.

Example (IOM)

```
MX9116N-A1# show smartfabric cluster

----------------------------------------------------------
CLUSTER DOMAIN ID : 119
VIP               : fde1:53ba:e9a0:de14:0:5eff:fe00:1119
ROLE              : BACKUP
SERVICE-TAG       : 3GB1XC2
MASTER-IPV4       : 10.11.105.15
----------------------------------------------------------
```

Example (VxRail - L2 fabric)

```
OS10# show smartfabric cluster

----------------------------------------------------------
CLUSTER DOMAIN ID : 100
VIP               : fde2:53ba:e9a0:cccc:0:5eff:fe00:1100
ROLE              : MASTER
SERVICE-TAG       : B37HXC2
MASTER-IPV4       : 10.11.106.27
PREFERRED-MASTER  :
----------------------------------------------------------
```

Example (VxRail - L3 fabric)

```
OS10# show smartfabric cluster

----------------------------------------------------------
CLUSTER DOMAIN ID : 100
VIP               : fde2:53ba:e9a0:cccc:0:5eff:fe00:1100
ROLE              : MASTER
SERVICE-TAG       : B37HXC2
MASTER-IPV4       : 10.11.106.27
PREFERRED-MASTER  : true
----------------------------------------------------------
```

Supported Releases

- MX9116n and MX5108n—10.5.0.1 or later
- SFS-supported OS10 switches—10.5.0.3 or later

show smartfabric cluster member

Displays information about the switch in a cluster. Information includes service tag, IP address, status, role, switch type and chassis model, and chassis service tag where the switch is connected.

Syntax `show smartfabric cluster member`

Parameters None

Default None

Command Mode EXEC

Usage Information When you run this command on a switch, the output that is displayed varies depending on the switch role. For example, if you run this command on the master switch, the output shows both the master and backup switch information. If you run this command on a backup switch, the output shows only the master switch information.

This command is supported in both Full Switch and SFS modes.

Example (IOM)

```
MX9116N-A1# show smartfabric cluster member
Service-tag    IP Address                             Status
  Role        Type    Chassis-Service-Tag Chassis-Slot
-------------------------------------------------------------
9GB1XC3        fde1:53ba:e9a0:de14:e6f0:4ff:fe3e:45dd  ONLINE
  MASTER      MX9116n SKY002L             B1
```

Example (VxRail)

```
OS10# show smartfabric cluster member
Service-tag IP Address                              Status
      Role   Type Chassis-Service-Tag Chassis-Slot
-----------------------------------------------------------
3Z4ZZP2     fde2:53ba:e9a0:cccc:54bf:64ff:fee6:e462 ONLINE
      BACKUP
3Z4ZZP1     fde2:53ba:e9a0:cccc:54bf:64ff:fee6:e463 ONLINE
      BACKUP
BR2ZZP2     fde2:53ba:e9a0:cccc:3c2c:30ff:fe49:2585 ONLINE
      BACKUP
B37HXC2     fde2:53ba:e9a0:cccc:e4f0:4ff:feb6:fdc3  ONLINE
      MASTER
G17HXC2     fde2:53ba:e9a0:cccc:e4f0:4ff:feb6:e1c3  ONLINE
      BACKUP
```

Supported Releases

- MX9116n and MX5108n—10.5.0.1 or later
- SFS-supported OS10 switches—10.5.0.3 or later

show smartfabric configured-server

Displays list of all configured servers information in a fabric. Information includes ID, model type, slot, chassis model and service tag, bonding technology, list of existing bond members, and server status.

Syntax `show smartfabric configured-server`

Parameters None

Default None

Command Mode EXEC

Usage Information This command is supported in both Full Switch and SFS modes.

Example (IOM)

```
MX9116N-B1# show smartfabric configured-server
----------------------------------------------------------
Service-Tag : 00FWX20
Server-Model : PowerEdge MX740c
Chassis-Slot : 1
Chassis-Model : POWEREDGE MX7000
Chassis-Service-Tag : SKY002L
Is-Discovered : TRUE
Is-Onboarded : TRUE
Is-Configured : TRUE
**********************************************************
Bonding Technology : LACP
BondMembers:
Nic-Id : Switch-Interface
----------------------------------------------------------
NIC.Mezzanine.1A-1-1 3GB1XC2:ethernet1/1/1
NIC.Mezzanine.1A-2-1 9A2HEM3:port-channel1
----------------------------------------------------------
```

Supported Releases
10.5.1.0 or later

show smartfabric configured-server configured-server-interface

Displays interface-level information of the configured servers. Information includes server ID, port ID, onboarded interface, server status, fabric ID, native VLAN, network profiles, and bandwidth partition details.

Syntax
`show smartfabric configured-server configured-server-interface` *`server-id`*

Parameters
`server-id`—Enter a configured server ID information.

Default
None

Command Mode
EXEC

Usage Information
This command is supported in both Full Switch and SmartFabric modes.

Example

```
MX5108N-B1# show smartfabric configured-server configured-server-
interface 004YX20 | no-more
----------------------------------------------------------
Server-Id             : 004YX20
----------------------------------------------------------
Port-Id               : NIC.Mezzanine.1B-2-1
Onboard-Interface     :
Fabric-id             :
Is-Discovered         : FALSE
Is-Onboarded          : FALSE
Is-Configured         : TRUE
NicBonded             : FALSE
Native-vlan           : 0
Networks              : c56d6202-0ec1-4fcd-b119-6abc761a1268
----------------------------------------------------------
Port-Id               : NIC.Mezzanine.1A-2-1
Onboard-Interface     : 1G86XC2:ethernet1/1/3
Fabric-id             :
Is-Discovered         : TRUE
Is-Onboarded          : FALSE
Is-Configured         : TRUE
NicBonded             : FALSE
Native-vlan           : 0
Networks              : c56d6202-0ec1-4fcd-b119-6abc761a1268
----------------------------------------------------------
Port-Id               : NIC.Mezzanine.1B-1-1
Onboard-Interface     : 2J86XC2:ethernet1/1/3
Fabric-id             :
Is-Discovered         : TRUE
Is-Onboarded          : FALSE
Is-Configured         : TRUE
NicBonded             : FALSE
Native-vlan           : 0
Networks              : c56d6202-0ec1-4fcd-b119-6abc761a1268
----------------------------------------------------------
Port-Id               : NIC.Mezzanine.1A-1-1
Onboard-Interface     :
Fabric-id             :
Is-Discovered         : FALSE
Is-Onboarded          : FALSE
Is-Configured         : TRUE
NicBonded             : FALSE
Native-vlan           : 0
Networks              : c56d6202-0ec1-4fcd-b119-6abc761a1268
```

Supported Releases
10.5.1.0 or later

show smartfabric details

Displays all details specific to the fabric. Details include name, description, ID, nodes that are part of the fabric, design type associated with the fabric, and status detail of a fabric.

Syntax `show smartfabric details`

Parameters None

Default None

Command Mode EXEC

Usage Information This command is supported in both Full Switch and SFS modes.

Example (IOM)

```
MX9116N-A1# show smartfabric details
-----------------------------------------------------------
Name               : A1-A2
Description        :
ID                 : fc6c9051-f499-4816-a54a-25ef6fef2e33
DesignType         : 2xMX9116n_Fabric_Switching_Engines_in_same_chassis
Validation Status: VALID
VLTi Status        : VALID
Placement Status : VALID
Nodes              : 3GB1XC2, 9A2HEM3
-----------------------------------------------------------
```

Example (VxRail)

```
OS10# show smartfabric details
-----------------------------------------------------------
Name               : AutoFab-08ee685b-d6d6-5d0c-99d2-ae78f800d4b7
Description        : Auto-Fabric Generator
ID                 : 08ee685b-d6d6-5d0c-99d2-ae78f800d4b7
DesignType         : AutoFabricDesign--1
Validation Status: VALID
VLTi Status        : VALID
Placement Status : VALID
Nodes              : CAC00N2, AZY1234
-----------------------------------------------------------
-----------------------------------------------------------
Name               : AutoFab-100
Description        : Auto-Fabric Generator
ID                 : 100
DesignType         :
Validation Status: VALID
VLTi Status        : VALID
Placement Status : VALID
Nodes              : AZY1234, CAC00N2, 9GTWNK2, FHTWNK2
-----------------------------------------------------------
```

Supported Releases

- MX9116n and MX5108n—10.5.0.1 or later
- SFS-supported OS10 switches—10.5.0.3 or later

show smartfabric discovered-server

Displays information about all the discovered servers. Information includes server tag, model, slot, chassis model, and chassis service tag.

Syntax `show smartfabric discovered-server`

Parameters None

Default None

Command Mode EXEC

Usage Information

This command is supported in both Full Switch and SmartFabric modes.

Example

```
MX5108N-B1# show smartfabric discovered-server
----------------------------------------------------------
Server-Tag               : 004YX20
Server-Model           : PowerEdge MX740c
Server-Slot            : 1
Chassis-Model          : POWEREDGE MX7000
Chassis-Service-Tag    : SKY002R
----------------------------------------------------------
----------------------------------------------------------
Server-Tag               : 0002X20
Server-Model           : PowerEdge MX740c
Server-Slot            : 1
Chassis-Model          : POWEREDGE MX7000
Chassis-Service-Tag    : SKY0044
----------------------------------------------------------
```

Supported Releases

10.5.1.0 or later

show smartfabric discovered-server discovered-server-interface

Displays interface-level information of all the discovered servers. Information includes port ID and switch interfaces on which the server is onboarded.

Syntax

`show smartfabric discovered-server discovered-server-interface server-id`

Parameters

`server-id`—Enter a discovered server ID information.

Default

None

Command Mode

EXEC

Usage Information

This command is supported in both Full Switch and SmartFabric modes.

Example

```
MX9116N-B1# show smartfabric discovered-server discovered-server-
interface 00FWX
20
Nic-Id : Switch-Interface
------------------------------------------------------
NIC.Mezzanine.1A-1-1 3GB1XC2:ethernet1/1/1
NIC.Mezzanine.1A-2-1 9A2HEM3:ethernet1/1/1
```

Supported Releases

10.5.1.0 or later

show smartfabric networks

Displays detailed description of the configured network profiles. Description includes network name, type, network ID, QoS priority type, and VLAN.

Syntax

`show smartfabric networks`

Parameters

None

Default

None

Command Mode

EXEC

Usage Information This command is supported in both Full Switch and SmartFabric modes.

Example (IOM)

```
MX9116N-A1# show smartfabric networks
Name           Type             QosPriority
    NetworkId                             Vlan
---------------------------------------------------
v5             GENERAL_PURPOSE  BRONZE
    8f018a8c-c355-4d81-9bee-85cfedcf8d2a  5
network100-105 GENERAL_PURPOSE  BRONZE
    deb0886c-4a9b-47f2-8220-55afcb1f1756  100 - 105
fcor           STORAGE_FCOE     PLATINUM
    d1de8f16-ebd0-4b1a-9689-a802d23b2b26  777
VLAN 1         GENERAL_PURPOSE  SILVER
    4bb446a3-702c-4a0f-abdd-07dd0c14775a  1
v1             GENERAL_PURPOSE  BRONZE
    9f2bed94-9148-46d8-9df6-3b606c83a472  500
```

Example (VxRail)

```
OS10# show smartfabric networks
Name                      Type   QosPriority
       NetworkId                 Vlan
--------------------------------------------
Client_Control_Network    VXLAN  IRON
       Client_Control_Network    3939
Client_Management_Network VXLAN  IRON
       Client_Management_Network 4091
```

Supported Releases

- MX9116n and MX5108n—10.5.0.1 or later
- SFS-supported OS10 switches—10.5.0.3 or later

show smartfabric nodes

Displays detailed inventory information about all the nodes in the fabric. Information includes service tag, type, status, mode, fabric ID associated with the node, chassis service-tag, and chassis-slot.

Syntax `show smartfabric nodes node-id` *`node-id`*

Parameters `node-id` *`node-id`*—Specify the service tag of a switch to view detailed information of that switch.

Default None

Command Mode EXEC

Usage Information This command is supported in both Full Switch and SmartFabric modes.

Example (IOM)

```
MX9116N-A1# show smartfabric nodes
Service-Tag Type    Status  Mode
   Chassis-Service-Tag        Chassis-Slot FabricId
---------------------------------------------------
3GB1XC2     MX9116n ONLINE  FABRIC
   SKY002L             A1
9GB1XC3     MX9116n ONLINE  FULL-SWITCH
   SKY002L             B1
9A2HEM3     MX9116n ONLINE  FABRIC
   SKY002L             A2
```

Example (VxRail)

```
OS10# show smartfabric nodes
Service-Tag Type      Status  Mode
      Chassis-Service-Tag  Chassis-Slot  FabricId
-------------------------------------------------
GGVQG02     S5232F-ON ONLINE  FABRIC
                                         7222c224-223c-5fa4-a244-
cd3ca1685550 (Name-Rack)
AZY1234     S5232F-ON ONLINE  FABRIC
```

```
                                                      7222c224-223c-5fa4-a244-
cd3ca1685550 (Name-Rack)
```

Example (VxRail)

```
OS10# show smartfabric nodes node-id GGVQG02
----------------------------------------------------------
Node Name            : Name-Leaf-2
Node Id              : GGVQG02
Node Type            : S5232F-ON
Node Status          : ONLINE
Node Mode            : FABRIC
Node Ready           : true
Node Model           : S5232F-ON
Replacement Node Id  :
Chassis service tag  :
Chassis slot         :
Fabric               : 7222c224-223c-5fa4-a244-cd3ca1685550 (Name-Rack)
Fabric node status   : OPERATIONAL
Software Version     : 10.5.2.xMRDEV
Hardware Version     : X01
Serial Number        : CN01WJVTCES0085G0037
----------------------------------------------------------
```

Supported Releases

- MX9116n and MX5108n—10.5.0.1 or later
- SFS-supported OS10 switches—10.5.0.3 or later

show smartfabric personality

Displays the personality of the SFS cluster.

Syntax `show smartfabric personality`

Parameters None

Default None

Command Mode EXEC

Usage Information

Use this command to identify the fabric personality of the SFS cluster. The output varies depending on the mode and roles of the switch, and personality. When this command is run on a switch, the system displays the personality and the roles of the switch. In VxRail deployment, if `vxrail` is displayed as personality it means a L2 fabric.

This command is supported in both Full Switch and SmartFabric modes.

Example (IOM)

```
MX9116N-A1# show smartfabric personality

Personality     :None
Role            :
ICL             :
```

Example (VxRail)

Full Switch mode:

```
OS10# show smartfabric personality

Personality     :None
Role            :
ICL             :
```

SmartFabric Services mode:

```
TOR1# show smartfabric personality

Personality     :vxrail
```

```
Role              :
ICL               :ethernet1/1/29, ethernet1/1/30
```

```
OS10# show smartfabric personality

Personality       :L3 Fabric
Role              :LEAF
ICL               :ethernet1/1/5, ethernet1/1/6
Leaf1#
```

```
OS10# show smartfabric personality

Personality       :L3 Fabric
Role              :SPINE
ICL               :
```

Supported Releases

- MX9116n and MX5108n—10.5.0.1 or later
- SFS-supported OS10 switches—10.5.0.3 or later

show smartfabric uplinks

Displays all uplink-related information in the SFS. Information includes uplink name, description, ID, media type, native VLAN, configured interfaces, and the network profile associated with the uplink.

Syntax `show smartfabric uplinks`

Parameters None

Default None

Command Mode EXEC

Usage Information This command is supported both in Full Switch and SmartFabric modes.

Example (IOM)

```
MX9116N-A1# show smartfabric uplinks
----------------------------------------------------------
Name                     : uplink to b1
Description              :
ID                       : 2725707d-886a-41c6-9d0d-38c4115788ff
Media Type               : ETHERNET
Native Vlan              : 1
Untagged-network         :
Networks                 : deb0886c-4a9b-47f2-8220-55afcb1f1756,
                           9f2bed94-9148-46d8-9df6-3b606c83a472
Configured-Interfaces : 9A2HEM3:ethernet1/1/42, 3GB1XC2:ethernet1/1/42
----------------------------------------------------------
----------------------------------------------------------
Name                     : u1
Description              :
ID                       : e1c8169e-00dd-4a72-9e42-54485c049591
Media Type               : FC
Native Vlan              : 0
Untagged-network         :
Networks                 : d1de8f16-ebd0-4b1a-9689-a802d23b2b26
Configured-Interfaces : 3GB1XC2:fibrechannel1/1/44:1
----------------------------------------------------------
```

Example (VxRail)

```
OS10# show smartfabric uplinks
----------------------------------------------------------
Name                     : FABRICUPLINKNew
Description              : L3VxLAN780 Uplink
ID                       : L3VxLANUplink-780
Media Type               : ETHERNET
Native Vlan              : 0
```

```
Untagged-network       :
Networks               : Network780
Configured-Interfaces : CAC00N2:ethernet1/1/22:2
----------------------------------------------------------
----------------------------------------------------------
Name                   : L3VLANUplink
Description            : Uplink On L3VLAN Network800
ID                     : L3VLANUplink-800
Media Type             : ETHERNET
Native Vlan            : 0
Untagged-network       :
Networks               : Network800
Configured-Interfaces : CAC00N2:ethernet1/1/22:1
----------------------------------------------------------
----------------------------------------------------------
Name                   : L2VxLANUplink
Description            : Uplink On L2VxLAN Network770
ID                     : L2VxLANUplink-770
Media Type             : ETHERNET
Native Vlan            : 0
Untagged-network       :
Networks               : Network770
Configured-Interfaces : CAC00N2:ethernet1/1/22:3
----------------------------------------------------------
----------------------------------------------------------
Name                   : L2VxLANUplink
Description            : Uplink On L2VxLAN Network771
ID                     : L2VxLANUplink-771
Media Type             : ETHERNET
Native Vlan            : 0
Untagged-network       : Network771
Networks               :
Configured-Interfaces : CAC00N2:ethernet1/1/22:4
----------------------------------------------------------
----------------------------------------------------------
Name                   : L3Uplink
Description            : Uplink On L3 Network600
ID                     : L3RUplink-600
Media Type             : ETHERNET
Native Vlan            : 0
Untagged-network       : Network600
Networks               :
Configured-Interfaces : AZY1234:ethernet1/1/21:2
----------------------------------------------------------
```

Supported Releases

- MX9116n and MX5108n—10.5.0.1 or later
- SFS-supported OS10 switches—10.5.0.3 or later

show smartfabric upgrade-status

Displays all the information about the upgrade status.

Syntax

```
show smartfabric upgrade-status
```

ⓘ **NOTE:** This command is accessible to users with `sysadmin`, `secadmin`, or `netadmin` roles.

Parameters None

Default Not applicable

Command Mode EXEC

Usage Information You can run this command on all the switches in the fabric. The output remains the same on all the switches.

Example

```
MX9116N-A1# show smartfabric upgrade-status

Opaque-id              : ea89b7c4-de00-4fc9-ad54-9ed5bf61e300
```

```
Upgrade Protocol      : PUSH
Upgrade start time    : 2021-02-12 01:51:29.261000
Status                : SUCCESS
Nodes to Upgrade      : SVC009F, SVC009A, 7H92Q03, SVC009C, 8PQXV23, 34HQXC2,
                        8PTXV23, D6RRNK2, BWGQXC2, 9KK1W23, 9J45W23, 8Q52W23
Reboot Sequence       : 7H92Q03,SVC009C,9KK1W23,SVC009F,8Q52W23,8PTXV23,D6RRNK2,
                        BWGQXC2,SVC009A,34HQXC2,5WQQXC2,9J45W23,8PQXV23

Node     Current      Curent    Status-Message
         -Action      -Status
--------------------------------------------------------------------------------
SVC009C  DOWNLOAD      SUCCESS  Skipping ONIE update for the node SVC009C since
                                ONIE is already installed with latest firmware.
SVC009F  DOWNLOAD      SUCCESS  Skipping ONIE update for the node SVC009F since
                                ONIE is already installed with latest firmware.
SVC009A  DOWNLOAD      SUCCESS  Skipping ONIE update for the node SVC009A since
                                ONIE is already installed with latest firmware.
8Q52W23  DOWNLOAD      SUCCESS  Skipping ONIE update for the node 8Q52W23 since
                                ONIE is already installed with latest firmware.
34HQXC2  DOWNLOAD      SUCCESS  Skipping ONIE update for the node 34HQXC2 since
                                ONIE is already installed with latest firmware.
BWGQXC2  DOWNLOAD      SUCCESS  Skipping ONIE update for the node BWGQXC2 since
                                ONIE is already installed with latest firmware.
9J45W23  DOWNLOAD      SUCCESS  Skipping ONIE update for the node 9J45W23 since
                                ONIE is already installed with latest firmware.
7H92Q03  DOWNLOAD      SUCCESS  Skipping ONIE update for the node 7H92Q03 since
                                ONIE is already installed with latest firmware.
9KK1W23  DOWNLOAD      SUCCESS  Skipping ONIE update for the node 9KK1W23 since
                                ONIE is already installed with latest firmware.
8PQXV23  DOWNLOAD      SUCCESS  Skipping ONIE update for the node 8PQXV23 since
                                ONIE is already installed with latest firmware.
D6RRNK2  ONIE-INSTALL  SUCCESS  [Action: REBOOT] Successfully rebooted.
8PTXV23  DOWNLOAD      SUCCESS  Skipping ONIE update for the node 8PTXV23 since
                                ONIE is already installed with latest firmware.
```

Supported Releases 10.5.2.4 or Later

show smartfabric validation-errors

Displays validation-error information of the topology. Information includes error category, subcategory, description, recommended action, severity, timestamp, EEMI, problem, and recommendation for each error.

Syntax `show smartfabric validation-errors`

Parameters None

Default None

Command Mode EXEC

Usage Information Use this command to view a list of topology validation errors with detailed description about each error. This command can be run on any IOM deployed in the same cluster.

This command is supported both in Full Switch and SmartFabric modes.

Example

```
OS10# show smartfabric validation-errors
----------------------------------------------------------
ErrorKey         : d77d0133-f8c8-4cd7-82b5-83266e5361eb-ISL-
[ICL-3_REVERSE]-NotFound-Issue
MessageID        :
Description      : Unable to validate the SmartFabric because
the VLTi cable for link ICL-3_REVERSE is not connected as per
fabric design 2xMX9116n_Fabric_Switching_Engines_in_same_chassis.
EEMI             : NFAB0012
Category         : FABRIC_ERROR
Subcategory      : ISL_ERROR
Severity         : SEVERITY_1
```

```
Recommended Action:Make sure that the VLTi cables are connected
to the correct ports as per the selected fabric design.
Timestamp        : 1587488570
Problem Link
    SourceNode          :HRA0028
    SourceInterface     :HRA0028:ethernet1/1/38
    DestinatioNode      :HRA0027
    DestinationInterface:
Recommended Link
    SourceNode          :HRA0028
    SourceInterface     :HRA0028:ethernet1/1/38
    DestinatioNode      :HRA0027
    DestinationInterface:HRA0027:ethernet1/1/38
----------------------------------------------------------

----------------------------------------------------------
ErrorKey         : 8ca24343-c819-4d0b-ab12-2b9d99a36079-ISL-
[ICL-2_REVERSE]-NotFound-Issue
MessageID        :
Description      : Unable to validate the SmartFabric because
the VLTi cable for link ICL-2_REVERSE is not connected as per
fabric design 2xMX5108n_Ethernet_Switches_in_same_chassis.
EEMI             : NFAB0012
Category         : FABRIC_ERROR
Subcategory      : ISL_ERROR
Severity         : SEVERITY_1
Recommended Action:Make sure that the VLTi cables are connected
to the correct ports as per the selected fabric design.
Timestamp        : 1587490907
Problem Link
    SourceNode          :HRA0038
    SourceInterface     :HRA0038:ethernet1/1/9
    DestinatioNode      :HRA0037
    DestinationInterface:
Recommended Link
    SourceNode          :HRA0038
    SourceInterface     :HRA0038:ethernet1/1/9
    DestinatioNode      :HRA0037
    DestinationInterface:HRA0037:ethernet1/1/9
----------------------------------------------------------
```

Supported Releases

- MX9116n and MX5108n—10.5.0.1 or later
- SFS-supported OS10 switches—10.5.0.3 or later

show switch-operating-mode

Displays the current operating mode of a switch.

Syntax `show switch-operating-mode`

Parameters None

Command Mode EXEC

Usage Information Some OS10 switches operate in both Full Switch and SFS modes. Use this command to view or verify the operating mode of a switch that is deployed in the SFS environment.

Example

```
OS10# show switch-operating-mode
Switch-Operating-Mode : SmartFabric Mode
```

Supported Releases 10.4.0E(R3S) or later

10

Best Practices

This chapter provides you a list of the best practices for SFS deployments with various solutions.

Recommended upgrade sequence for SFS and solution deployments

This section covers the recommended upgrade sequence for SFS and its supported solution deployments. This recommended order of upgrade is applicable for all SFS deployments.

Prerequisites

Ensure that you follow the prerequisites before proceeding with the upgrade.

- Check the release notes of each product for the version you are upgrading to for any caveats or upgrade sequence requirements.
- Verify if all the product version that you are upgrading to are compatible with the solutions support matrix, see Dell EMC Networking Solutions Support Matrix.

Upgrade sequence in deployments with OMNI

This recommended order of upgrade is applicable for all SFS deployments with OMNI.

1. Upgrade OMNI. Follow the instructions provided in the OMNI User Guide and the OMNI Release Notes.
2. Upgrade OS10 using OMNI.
   a. Before upgrading OS10, check the OS10 release notes of the respective OS10 version and the *Dell EMC SmartFabric OS10 Installation, Upgrade, and Downgrade Guide* available in SmartFabric OS documentation for any additional requirements.
   b. Follow the SmartFabric OS upgrade procedure in the OMNI User Guide.

   > **NOTE:** The switches in the SmartFabric must be at the same version of OS10 when upgrading using OMNI. If not, upgrade process may fail. If you want to upgrade the switches that have different OS10 versions installed, upgrade each switch manually.
3. If needed, upgrade the solutions deployed with SFS such as VxRail, ESXi, and so on. See the respective solution documentation for upgrade instructions.

Upgrade sequence in deployments without OMNI

This recommended order of upgrade is applicable for all SFS deployments that do not include OMNI.

1. Upgrade OS10 on the switches, see *Dell EMC SmartFabric OS10 Installation, Upgrade, and Downgrade Guide* available on the SmartFabric OS10 documentation support page. Check the OS10 release notes of the respective OS10 version for any additional requirements.
2. If needed, upgrade the solutions deployed with SFS such as VxRail, ESXi, and so on. See the respective solution documentation for upgrade instructions.

Frequently asked questions

This chapter provides answers to frequently asked questions about SFS.

- **Can I use different platforms in the same VLT leaf pair in a fabric?**

No, both leaf switches in the same VLT pair must be the same platforms in SFS deployments. See Networking Solutions Support Matrix for the list of supported switches and typical roles in SFS deployments.

- **Can I use the different switch port profiles available on S4148F-ON in SmartFabric mode?**

  Yes, all six profiles available on S4148F-ON can be used in SmartFabric mode as follows:

  1. Set the appropriate port profile for the switch in Full-Switch mode.

     CAUTION: Changing the switch port profile resets the system defaults and overwrites the existing switch configuration. You must reconfigure the setting on a switch manually after applying the new port profile.
  2. Reload the switch to apply the new port profile.
  3. Enable SFS on the switch to change the switch mode to SmartFabric. After enabling SFS, the switch reboots in SmartFabric mode and uses the set port profile.

  NOTE: For more information about port profiles, see *Dell EMC SmartFabric OS10 User Guide*. This procedure is applicable only on S4148F-ON platform as S4148FE-ON and S4148U-ON switches are not supported in SFS deployments.

- **How do I change the OS10 admin password in SmartFabric mode?**

  By default, `admin` is the username and password to log in to the OS10 CLI. The default OS10 admin password should be changed on each switch after the first login. The procedure is the same for a switch as in Full Switch mode or SmartFabric mode. For more information about changing the password, see *Dell EMC SmartFabric OS10 User Guide*. The admin password does not have to be the same password on all switches in a SFS deployment. When logging in to the SFS GUI, so use the credentials of the leader switch. The leader switch may change and use the command `show smartfabric cluster` at any switch in the SmartFabric to see the details of the leader switch.

- **How can I upgrade firmware for a SFS-enabled switch?**

  The firmware can be upgraded using the `image install` command. For more information about installing firmware upgrade, see *Dell EMC SmartFabric OS10 User Guide*.

- **Can I have a peer switch with multiple BGP devices?**

  Yes, you can peer SFS switches with multiple BGP devices.

- **Is OSPF supported for SFS?**

  No. OSPF is not supported in SFS with leaf and spine deployments.

- **How many jump ports can I configure?**

  You can configure one jump port per leaf switch in a fabric.

- **Can I change the VLT interface after SFS has been enabled?**

  Yes. There is an option to change the VLTi interfaces in SFS mode using the `smartfabric vlti` CLI command, see smartfabric vlti.

11

Appendix

This chapter covers additional information that can support you with the fabric configuration tasks.

REST_USER account

When the fabric is formed, the SFS REST service is started on the leader node. Applications that are integrated with SFS use the REST service for fabric operations. Communication is performed with the fabric using the IPv6 VIP assigned to the SFS leader, or using the IPv4 out-of-band Management IP of the leader.

A default REST_USER account is created to authenticate all REST queries. OMNI communicates with SmartFabric REST Services through REST_USER account only. The default password is `admin`, and Dell Technologies recommends changing the password through VxRail Manager or OMNI. For more information about how to change the REST_USER account password from OMNI, see *OpenManage Network Integration for SmartFabric Services User guide*.

Internal fabric components and networks

SFS automatically creates the following components and networks:

Internal SFS components

SFS creates a VLT fabric automatically in the leaf and spine environment. VLT fabric is autoassigned with a fabric-ID, a universally unique identifier (UUID). When a VLT fabric is created, the management IP addresses of the VLT peers are used automatically to set up the VLT backup link. If the management IP address of the peers is changed after the fabric is created, the VLT backup link is updated automatically.

SFS creates a network fabric with the leaf and spine switches automatically. The network fabric is autoassigned with a fabric-ID, a name, and description. The fabric name and description are automatically assigned, but can be changed through the SFS GUI. For more information, see Update fabric name and description.

Fabric links create a connection between the switches in a network fabric. ISL is a link formed between a leaf and spine switch. All parallel links with same connectivity are grouped to form a LAG interface. ICL or VLT interconnect (VLTi) is a link formed between the two leaf switches in a same rack.

Internal virtual networks

The internal virtual networks created by SFS are:

VLAN 4000 for SFS cluster control
SFS automatically configures VLAN 4000 on all the switches that are discovered in a fabric, and uses the network for all internal fabric operations. When a leaf or spine switch is discovered, the ICL or ISL ports are automatically added as tagged members.

VLAN 4001 to 4079 for leaf and spine connections
SFS automatically configures the leaf and spine network using eBGP as the underlay routing protocol. SFS uses the reserved VLAN range from 4001 to 4079 with automatic IP addressing to set up the peer connections. When SFS detects an ISL connection on either a leaf or spine switch, it assigns the ISL to the untagged member of this VLAN. IP address from reserved range is used for this VLAN, and an eBGP session is started on the VLAN interface.

VLAN 4080 for Global untagged VXLAN
SFS automatically configures VXLAN overlay networks with EVPN to extend networks between racks in a multirack deployment. VLAN 4080 with automatic IP addresses from the reserved range is used for ICL links. VXLAN requires one VLAN to be assigned globally for untagged port-scoped VLAN (Port, VLAN) pairs.

VLAN 4089—OS10 internal use

In SmartFabric mode, VLAN 4089 is the default VLAN and is reserved for OS10 internal use.

VLAN 4090—iBGP peering between leaf switches

SFS automatically configures iBGP peering between a pair of leaf switches directly connected over ICL links. VLAN 4090 is created automatically with IP addresses from reserved range for enabling iBGP sessions between the VLT peer switches.

VLAN 4094—VLT control VLAN

SFS automatically creates VLAN 4094 on all leaf switches. This VLAN is used for all VLT control traffic between two VLT peer switches and is added on the VLT ICL ports on leaf switches.

VLAN 4091—Default client management network

SFS automatically configures an overlay network that is called a `client_Management_Network`. When a device is automatically onboarded on to the network fabric, the device uses the VLAN mapped to this overlay network. This network is a native VLAN unless there is a policy specifying a different native VLAN. VLAN 4091 is used as the default client management VLAN for the VXLAN network.

NOTE: You can change this VLAN to a specified VLAN through SFS GUI.

VLAN 3939—Default client control network

SFS configures a second overlay network that is called `Client_Control_Network` for SFS-integrated solutions. When a device such as VxRail is discovered, it is automatically added as a tagged member of this network. SFS enables leader advertisement and fabric discovery by integrated solutions. The SFS leader virtual IP address for VXLAN network is advertised. The VIP address `fde1:53ba:e9a0:cccc:0:5eff:fe00:1100` is fixed and not user configurable.

VLAN 3939 is used as the default client control VLAN for this VXLAN network for SFS-integrated solutions including VxRail and PowerScale solutions. Although you can change the VLAN associated with the default client management and control networks, Dell Technologies recommends not to change the VLANs for VxRail deployments.

NOTE: For more information about other internal general networks created by SFS, see Networks.

To check the networks that are created in SFS-integrated deployment, use `show virtual-network` command. Following is the example output :

```
OS10# show virtual-network
Codes: DP - MAC-learn Dataplane, CP - MAC-learn Controlplane, UUD - Unknown-Unicast-Drop
Un-tagged VLAN: 4080
Virtual Network: 3939
   Description: In-band SmartFabric Services discovery network
   VLTi-VLAN: 3939
   Members:
      VLAN 3939: port-channel1000, ethernet1/1/12, ethernet1/1/13
   VxLAN Virtual Network Identifier: 3939
      Source Interface: loopback2(172.30.0.0)
      Remote-VTEPs (flood-list): 172.30.0.1(CP)

Virtual Network: 4091
   Description: Default untagged network for client onboarding
   VLTi-VLAN: 4091
   Members:
      Untagged: ethernet1/1/12, ethernet1/1/13
      VLAN 4091: port-channel1000
   VxLAN Virtual Network Identifier: 4091
      Source Interface: loopback2(172.30.0.0)
      Remote-VTEPs (flood-list): 172.30.0.1(CP)
```

Reserved networks

SFS uses the 172.16.0.0/16 and 172.30.0.0/16 networks internally for the leaf and spine network configuration. If these networks conflict with any networks in the existing deployment, change the default networks using the instructions that are provided in Edit Default Fabric Settings.

Domain ID and ASN mapping

d

After you enable the SFS on the switches, SFS:

- Associates domain ID for the cluster if it is not assigned when enabling SFS. By default, the domain ID is set to 100.
- Autogenerates ASN for the leaf and spine switches in the fabric.
- Autogenerates subnet with prefix from domain ID.

ⓘ **NOTE:** You can edit the domain ID from Fabric Settings page. Any change to the domain ID changes the ASN for leaf and spine switches, and the subnet details.

The following table lists the ASN and BGP subnet and VTEP subnet mapped with the domain ID:

Table 9. Domain ID to ASN mapping

Domain ID	Leaf ASN	Spine ASN	Fabric BGP subnet and mask	Fabric VTEP subnet and mask
100	65011	65012	172.16.0.0/16	172.30.0.0/16
101	65015	65016	172.18.0.0/16	172.32.0.0/16
102	65017	65018	172.19.0.0/16	172.33.0.0/16
103	65019	65020	172.20.0.0/16	172.34.0.0/16
104	65021	65022	172.21.0.0/16	172.35.0.0/16
105	65023	65024	172.22.0.0/16	172.36.0.0/16
106	65025	65026	172.23.0.0/16	172.37.0.0/16
107	65027	65028	172.24.0.0/16	172.38.0.0/16

MSTP Support on L3 personality

The default spanning tree mode in SFS is RPVST+. On Dell EMC PowerSwitches in SFS mode, RPVST+ is enabled globally and automatically configured.

In L3 personality, STP is enabled on:

- SFS cluster control VLAN (VLAN 4000). The spine switches are configured to take over the STP root role.
- All user created VLANs.

In L3 personality, STP is disabled on:

- All inter leaf-spine VLANs and leaf-leaf VLANs (4001-4091).
- All server facing ports.

ⓘ **NOTE:** Do not modify STP settings on switches in SFS L3 personality.

If you must interoperate switches that are controlled by SFS to external switches which are running RSTP or MSTP, SFS has an API to change the global fabric STP mode to MSTP. SFS creates 2 reserved MSTIs:

- MST with instance id 63: Cluster control VLAN 4000 is part of this instance. Spine switches are configured to take over as STP root for this MSTI.
- MST with instance id 62: All the SFS reserved VLANs (4001-4091) are part of this configuration. On this MSTI, spanning tree is disabled.
- All user created VLANs are part of CST (default MST instance), which interoperates with RSTP. STP is enabled for this MSTI.

You can change the mode to MSTP once the fabric is created. When you change the mode, the whole fabric goes through a reboot cycle and the new mode is set to MSTP.

ⓘ **NOTE:** Changing the STP mode impact the traffic flow in the cluster.

When the mode is changed, MSTI is created and VLANS are assigned to the MSTI. The CST is configured with STP priority such that SFS-controlled switches have lower priority to become a root bridge.

There is no change on existing STP behavior for SFS-controlled entities because of this change. All other STP behaviors such as disabling STP on server facing ports still holds good.

You can enable MSTP or revert to RPVST using the **Edit Default Fabric Settings** option available in the SFS UI. Select **STP Mode** as MST or Rapid PVST For more information, see Edit Default Fabric Settings.

ⓘ **NOTE:** The default spanning-tree priority value that is configured on a SFS-enabled switch for VLANs or instance 0 is: default priority for RPVST is 32769 and instance 0 priority for MSTP is 61440.

Networks

This section describes the following type of networks that you can create and associate these networks with the entities present in a fabric:

- General purpose networks
- VXLAN networks
- L3 VLAN networks
- L3 Routed networks
- Multirack L3 VLAN network

  ⓘ **NOTE:** Prior to 10.5.3.0 release, you cannot configure this network using SFS GUI. This network template is applicable for NSX-T deployment and you can configure this network using OMNI UI. For more information about this network type, see *OpenManage Network Integration for SmartFabric Services User Guide, Release 2.0*.

You can create these networks using SFS GUI and then associated these networks with servers profiles, uplinks, or interfaces for traffic flow. For more information, see Create a network.

General purpose networks

General purpose networks can be categorized as L2 VLAN networks in SFS L2 personality and L2 VXLAN networks in SFS L3 personality. In L3 personality, when you create a general purpose network, SFS automatically creates a virtual network (VXLAN) corresponding to a VLAN network. This virtual network has one-to-one mapping with the network, for each VLAN there exists a virtual network with VNI same as the VLAN ID. If you delete a VLAN network, it automatically deletes the associated VXLAN network.

For example, if you create a L2 general purpose network with VLAN ID 50, SFS creates a VXLAN network with VNI 50 and associates to VLAN 50.

VXLAN networks

VXLAN network extends L2 connectivity over an underlay L3 connected network. Association of VXLAN network to interface creates a binding and associates this interface to VXLAN bridge. L3 VXLAN network supports asymmetric-IRB. Create a virtual network template and a network template, and associate the virtual network template to network template.

- The virtual network template defines the VNET-ID.
- Network template defines the VLAN ID.

L3 VXLAN network is a VXLAN type of network that contains a list of IP addresses and an anycast IP address. Optionally, you can specify DHCP relay addresses. L3 VXLAN network can be configured over a leaf switch. L3 VXLAN network can be attached to an uplink. Each VLTi uplink interface contains an IP address that is allocated from the list of IP addresses that are configured on the L3 VXLAN network.

L3 VLAN network

L3 VLAN network is used for L3 VLAN underlay. Specify:

- VLAN ID
- Pair of IP addresses to be assigned to the VLT pair
- VRRP gateway IP address for VIP

L3 VLAN network contains a list of IP addresses and a gateway IP address. Optionally, you can specify DHCP relay addresses. You can configure and attach a L3 VLAN network to an uplink. Each VLTi uplink interface contains an IP address that is allocated from the list of IP addresses that are configured on the L3 VLAN network.

ⓘ **NOTE:** You must assign a L3 VLAN only to the server interface profile configured in a single rack.

L3 Routed network

L3 routed network is used to assign IP address on a single interface. L3 routed network contains a list of IP addresses and a gateway IP address. Optionally, you can specify DHCP relay addresses. L3 routed network can be configured on a leaf or a spine switch. L3 routed network can be attached to an uplink. Each uplink interface contains an IP address that is allocated from the list of IP addresses that are configured on the L3 routed network. An LACP port channel cannot be the remote end for these uplinks. If gateway IP address is specified, then VRRP is enabled and the switches configure this IP address as the gateway IP address. Attach this network to any uplink that has a single interface. This network can only be attached to a single entity.

Multirack L3 VLAN network

Multirack L3 VLAN network is a template that contains rack-specific IP configurations for a L3 VLAN. You can configure IPv4 attributes for each rack in an NSX-T deployment consisting of multirack leaf and spine topology. A rack holds a pair of switches that are configured with VLT. The rack identifier is the same as the fabric ID. A rack containing a single switch without VLT does not support multirack L3 VLAN network.

In an NSX-T topology that is connected to an SFS cluster, you can create a multi rack L3 VLAN network. Specify the network identifier, L3 VLAN ID, and list of rack-specific IPv4 configurations for this network. Attach the created network as untagged or tagged on any uplink or server interface of the rack. After you attach the network, SFS automatically creates a L3 VLAN and applies the rack-specific IP configurations to the respective rack within the topology. You can modify the IPv4 configuration for a specific rack in the multi rack L3 VLAN network.

ⓘ **NOTE:** For faster traffic convergence, Dell Technologies recommends you to set the BGP keepalive timer to 1 second and hold-time to 3 seconds on the NSX-T tier-0 node when the BFD configuration is disabled or not supported. If BFD is enabled, BFD timer setting is used. For more information, see BFD support information in SFS configuration notes.

Uplinks

This information explains the types of uplinks that you can create in a fabric.

L2 uplinks

You can configure L2 uplinks only on the leaf switches in a fabric. L2 uplinks are a set of user-selected ports that belong to same VLT leaf switches in which the L2 network is applied. The L2 uplink from the leaf switches is either an LACP or a static LAG. SFS creates a VLT LAG for connected ports. If the ports are from a single device, the VLT LAG is a single armed VLT LAG. In case these ports exist on multiple switches, a VLT or LAG is formed across these ports. This LAG can be made as an access point for a VXLAN L2 network.

L3 uplinks

You can configure L3 uplinks on a leaf or spine switch. SFS supports L3 routed or L3 VLAN uplinks.

With L3 routed uplinks, point-to-point links are required between the switches with L3 uplinks and the external switches. With L3 VLAN, all uplinks are in a LAG, and an IP address is assigned to the VLAN containing the LAG. Point-to-point IP networks and addresses must be planned for each physical link in the L3 uplink. Each leaf switch in the fabric needs an IP address on the external Management VLAN, and an anycast gateway address on the same VLAN. The virtual router or anycast gateway address is shared by all leaf switches in the fabric.

Routing profiles

SFS supports eBGP and static routing profiles. You can create and delete the routing profiles, but you cannot edit any routing policy. To modify, you must delete the existing routing policy and create a policy with the relevant configuration.

Static route—A static route profile is a routing template that contains a network prefix and the next hop IP address. During uplink configuration, you can associate this profile to the uplinks created on one or more switches in the fabric. When this profile is created, a route with the specified prefix and next hop IP address is configured on the switch.

eBGP route—An eBGP peer routing profile is a routing template that contains BGP remote addresses and the remote AS number. A remote address can be an interface address or a loopback address. During uplink creation, you can associate this policy to uplink created on one or more switches in a fabric. When this policy is created, a BGP session is configured on the switch.

ⓘ **NOTE:** SFS does not support OSPF or routing protocols other than eBGP.

Uplink bonding options

Following are the uplink bonding types supported:

- LACP
- Static bonding

LACP

In LACP uplink bonding, SFS configure the LACP LAG for the uplink using the LACP PDUs received from the remote device. Networks that are attached to the uplink are associated with the LACP LAG that is created.

Static bonding

In static bonding, SFS configures a static LAG for the uplink and the networks that are attached on the uplink are associated with the LAG that is created.