OpenManage Network Integration for SmartFabric Services User Guide

Release 2.0

December 2020
Rev. A00

DELL EMC

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.

Contents

Chapter 1: Change history **5**

Chapter 2: Overview of OMNI, SFS, VxRail, and PowerEdge MX **7**

Chapter 3: SmartFabric Services **10**
SFS for leaf and spine fabrics 10
SFS initial setup 10
Enable SFS 11
Fabric creation 11
SFS and OMNI supported solutions 14
SFS personalities 15

Chapter 4: OpenManage Network Integration **17**
Install OMNI virtual appliance using vCenter 18
Upgrade OMNI appliance 24
Set up OMNI 27
Install OMNI application on ESXi server without vCenter 38
OMNI appliance console CLI menu 43
Generate and install SSL certificate 44
View and configure docker private network settings 47

Chapter 5: OMNI vCenter integration **51**

Chapter 6: Access the OMNI stand-alone portal **53**
Access the OMNI Fabric Management Portal 53
Configure OMNI 54
Register vCenter with OMNI 55
Access OMNI plug-in from the vCenter 60
Edit OMNI autodiscovered SmartFabric instance 60
Add SmartFabric instance 61
OMNI support for vCenter Enhanced Linked mode 65
Host network inventory 66
View service instance and vCenter relationships 69
OMNI Information 70
OMNI Appliance Management user interface 71

Chapter 7: OMNI automation support for PowerEdge MX SmartFabric **77**
Workflow to integrate OME-Modular with OMNI 77
OMNI VLAN automation 81

Chapter 8: OMNI automation support for NSX-T **83**
Workflow to integrate NSX-T with OMNI 83
OMNI automation for NSX-T 87

Chapter 9: OMNI support for SmartFabric instances.....90
OMNI feature support matrix.....91
View SmartFabric instance overview 91
View node details.....92
View fabric topology.....94
Manage switches in a fabric.....95
View switch and port details.....95
Edit port configuration on a switch.....96
Manage unused switch ports.....101
Configure breakout ports.....102
Add a jump port.....105
Configure server interface profile.....108
Create server interface profile.....108
Edit networks and ports in a server interface profile.....115
Delete a server interface profile.....117
Import ESXi host profiles from vCenter.....117
Import SmartFabric discovered server interfaces.....121
Configure and manage uplinks.....125
Create L2 Uplink.....126
Create L3 Uplink.....128
Edit networks and ports in an uplink.....133
Delete an uplink.....135
Configure networks and routing configuration.....135
Configure networks.....136
Configure Routes.....153
Configure global settings for SmartFabric.....159
View fabric events and compliance status.....166
View fabric events.....166
View fabric compliance status.....167

Chapter 10: Lifecycle management.....169
Change SmartFabric password.....169
Upgrade SmartFabric OS in switch.....169
Replace switch in a fabric.....172
Back up and restore the fabric configuration.....173
Restore from a backup file.....179

Chapter 11: Troubleshooting.....181
Logs and support data for troubleshooting.....181
Verify OMNI VM connectivity.....181
Unable to add SmartFabric instance in OMNI.....184
Missing networks on server interfaces.....184
Unable to launch OMNI UI.....185
OMNI plug-in does not show service instance.....186
Unable to register the vCenter in OMNI.....186
OMNI is unable to communicate with other devices.....187
Timestamp not synchronized in OMNI.....187
Reset OMNI VM password.....189

1

Change history

The following table provides an overview of the changes to this guide from a previous OMNI release to the OMNI 2.0 release. For more information about the new features, see the respective sections.

Table 1. New in 2.0

Revision	Date	Feature	Description
A00	2020–12-16	OMNI automation support for PowerEdge MX SmartFabric	OMNI manages fabric automation for ESXi hosts deployed within the Dell EMC PowerEdge MX solution.
		OMNI automation support for NSX-T	OMNI supports fabric automation for NSX-T Manager integration with SmartFabric Services.
		Register vCenter through OMNI Fabric Management UI	Register vCenter instance using OMNI Fabric Management UI.
		Install OMNI VM on ESXi server without vCenter	Deploy the OMNI appliance on a VMware ESXi server using the OMNI OVA file.
		Relationship information	View relationship between the vCenter and service instances (SmartFabric, NSX-T Manager, and OME-M instances).
		OMNI SmartFabric instance overview	OMNI displays the summary overview of key metrics such as device status and health, latest fabric events, and fabric compliance errors for the SmartFabric instance.
		OMNI Home page enhancement	OMNI Home enhancement with an option to add different service instance separately.
		Support for onboarding unknown server discovered interfaces	OMNI supports dynamic onboarding of unknown servers that are discovered by SmartFabric.
		Configuration support for SmartFabric global settings	Configure the SmartFabric switch services settings through OMNI Fabric Management UI: • NTP • DNS • Syslog • SNMP
		OMNI Secure sign on support	Secure sign on enhancement for OMNI.
		vCenter Enhancement Linked mode	OMNI support for vCenter Enhanced Linked mode.
		Fabric events	View latest fabric events for each SmartFabric instance.

Table 1. New in 2.0 (continued)

Revision	Date	Feature	Description
		Configure docker private network settings	View and configure docker private network settings on the OMNI appliance.
		Fabric compliance	View fabric compliance status and the recommended action for each SmartFabric instance.

2

Overview of OMNI, SFS, VxRail, and PowerEdge MX

Enterprises are adopting the power of automation to transform their IT operations, and enable a more agile and responsive infrastructure in their data center. Network operators must leverage the power of automation within and across their departmental functions, delivering integrated solutions which cater to cloud-based consumption models.

SmartFabric Services

SmartFabric Services (SFS) is an automation framework that is built into Dell EMC SmartFabric OS10, to integrate converged and hyperconverged infrastructure systems. These solutions deliver autonomous fabric deployment, expansion, and life cycle management.

SFS enables converged infrastructure (CI) and hyperconverged infrastructure (HCI) for system administrators to deploy and operate the network fabric for the infrastructure solution as an extension of the solution being deployed. This integrated network fabric is built using industry-standard protocols adhering to the best practice recommendations for that solution, and is interoperable with existing data center networks deployment.

There are two types of SFS:

1. SFS for Leaf and Spine – supported on selected Dell EMC S-series and Z-series PowerSwitches.
2. SFS for PowerEdge MX – supported on PowerEdge MX switches.

For more information regarding supported switches, see SmartFabric OS10 Support Matrix.

OpenManage Network Integration

Dell EMC OpenManage Network Integration (OMNI) is a management application that is designed to complement SFS, providing a web-based UI for operating one or more automated network fabrics deployed using SFS (called SmartFabric instances).

OMNI is delivered as a virtual appliance which can be deployed as:

- A stand-alone virtual machine enabling a web portal to manage one or more SmartFabric Instances
- An external plug-in for VMware vCenter. When deployed as a plug-in for VMware vCenter, OMNI enables:
  - Zero-touch automation of physical underlay network fabric running SFS corresponding to changes in the virtual network layer
  - Extension of vCenter Host Network Inventory data to include physical switch connectivity details for easy monitoring and troubleshooting
  - A single pane of management for one or more SmartFabric instances through the OMNI portal pages that are embedded within vCenter

VxRail SFS integration

Dell EMC VxRail integrated with SFS automates and simplifies networking for VxRail hyperconverged infrastructure deployments and ongoing network operations. As hyperconverged domains scale, the network fabric becomes the critical piece of successful deployment. VxRail integration with SFS allows customers to deploy network fabrics for VxRail clusters as an extension of the VxRail clusters without extensive networking knowledge. The network fabric is automatically configured for the VxRail nodes as the operators deploy their VxRail clusters.

Key benefits

- Faster time to production
  - Plug and play fabric formation for VxRail.
  - VxRail Manager automatically creates fabric policies for VxRail nodes.

- SmartFabric automates all fabric functions.
- Integrated life cycle
  - Fabric creation, expansion, and maintenance follow the VxRail application model.
  - HCI fabric operations are fully managed through VxRail Manager or vCenter.
- Better infrastructure visibility
  - Tight integration between VxRail appliance and Dell EMC ON-Series PowerSwitches.
- Improved SLA
  - Fully validated software stack.
  - Protection from human-error due to predictable and repeatable HCI fabric experience.
- Enhanced support experience
  - World-class Dell EMC HCI and fabric services.
  - Fabric that is integrated into VxRail services and support experience.

Required components

- Dell EMC PowerSwitches supporting SmartFabric Services.
- Dell EMC SmartFabric OS10 for PowerSwitch models.
- Dell EMC OpenManage Network Integration (OMNI).
- Dell EMC VxRail hyperconverged nodes when deploying VxRail integrated solution.
- VMware vCenter internal to VxRail cluster or existing vCenter in customer environment.

See the SmartFabric OS10 Support Matrix for the latest software releases that support the VxRail and SmartFabric Service integrated solution.

Supported switches

Table 2. Supported switches for VxRail-SFS

PowerSwitches	Switch type	VxRail node connectivity options
• S4112F-ON • S4112T-ON • S4128F-ON • S4128T-ON • S4148F-ON • S4148T-ON	Leaf or top of rack switches	10 GbE
• S5212F-ON • S5224F-ON • S5248F-ON • S5296F-ON		25 GbE
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	Spine	—

S4248FB-ON, S4248FBL-ON switches are supported for solutions without VxRail.

PowerEdge MX integration

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. As part of the PowerEdge MX platform, Dell EMC SmartFabric OS10 includes SmartFabric Services which is fully integrated with the MX platform.

In MX, a SmartFabric is a logical entity that consists of a collection of physical resources, such as servers and switches, and logical resources such as networks, templates, and uplinks. The OpenManage Enterprise - Modular (OME-M) console provides a single interface to manage these resources as a single unit.

Key benefits

- Data center modernization
  - I/O aggregation.
  - Plug and play fabric deployment.
  - Single interface to manage all switches in the fabric.
- Lifecycle management
  - SmartFabric OS10 updates across the fabric.
  - Automated or manual rollback to last well-known state.
- Fabric automation
  - Physical topology compliance.
  - Server networking managed using templates.
  - Automated QoS assignment per VLAN.
  - Automated storage networking.
- Failure remediation
  - Dynamically adjusts bandwidth across all interswitch links when there is a link failure.
  - Automatically detects fabric misconfigurations or link level failure conditions.
  - Automatically heals the fabric on failure condition removal.

When PowerEdge MX switches are in SmartFabric Services mode, they operate entirely as a Layer 2 network fabric. Layer 3 protocols are not supported. For more information about MX switches, see MX documentation.

More resources

List of more resources you may need:

Table 3. More resources

Path and Links to Documents	Description
Dell EMC Networking OS10 Info Hub > **OS10 User Guides** > **OS10 Dell EMC SmartFabric OS10 User Guide, 10.5.2**	This document contains information to help you understand, configure, and troubleshoot your OS10 networking operating system.
Dell Technologies VxRail Networking Infohub > **Guides**	This page contains reference documents to configuration, deployment, and other guides for VxRail networking solutions.
SmartFabric OS10 Solutions Support Matrix	This page contains the various support matrices of SmartFabric OS10 solutions including VxRail, PowerStore, Isilon front-end, PowerEdge ESXi, vSAN Ready Nodes, and PowerEdge MX.
Dell EMC OpenManage Network Integration for VMware vCenter > **Manuals and documents**	This page lists the OMNI reference manuals from previous versions.
Dell EMC PowerEdge MX Modular Switches	This page list the manuals and reference documents to configure PowerEdge MX Modular Switches in different deployment scenarios.

3

SmartFabric Services

SFS offers plug and play data center network fabric deployment, expansion, and management of Dell EMC infrastructure as turnkey solutions. SFS is a component of SmartFabric OS10 that provides the framework to automatically deploy the network as a single logical entity which enables the integration of Dell EMC infrastructure solutions.

SFS offers turnkey network solution for data center infrastructure using Dell EMC PowerSwitch and PowerEdge MX data center switches.

This information provides an overview of the SFS solution that is built on an automated data center leaf and spine network fabric using Dell EMC PowerSwitch models.

This section provides information about SFS for PowerSwitch infrastructure. For more information about SFS for PowerEdge MX, see Dell EMC PowerEdge MX Infohub.

SFS for leaf and spine fabrics

SFS is built on top of modern leaf and spine data center design that is optimized for the increased east-west traffic requirements of modern data center workloads. The entire leaf and spine network fabric is orchestrated and managed as a single object, eliminating the need for box-by-box configuration and management of the switches.

The fabric can start from a single rack deployment with two leaf or top-of-rack (ToR) switches, and expanded to a multi rack leaf and spine network fabric. The fabric is automatically built and expanded using industry-standard Layer 2 and Layer 3 protocols as new switches are connected.

L3 fabric profile

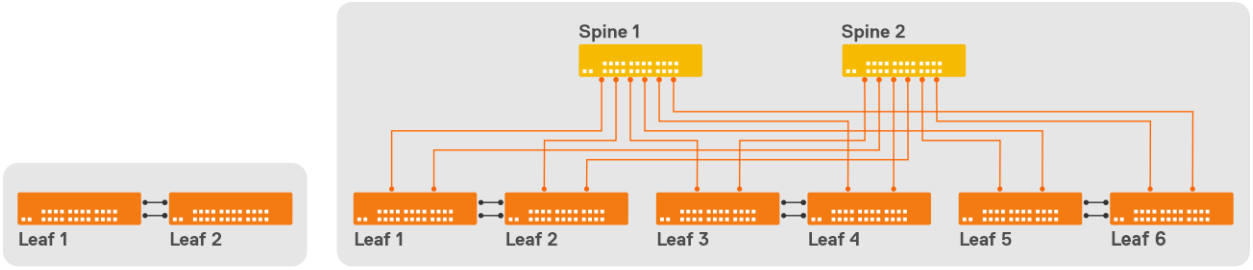

Single Rack Fabric

Multi Rack Fabric

NOTE: SmartFabric Services can be enabled when there are at least two leaf/ToR switches connected as a VLT pair.

SFS initial setup

When PowerSwitch models with SmartFabric OS10 power on, the switches are operating in the normal Full Switch mode. This information explains how to start the automated discovery and fabric creation process.

1. Log in to each switch console.
2. Configure the out-of-band Management IP address.
3. Upgrade SmartFabric OS10 to supported versions based on the Support Matrix.
4. Enable SmartFabric Services on the switches.

For complete information about configuring the out-of-band Management IP address and upgrading the switch operating system, see *Dell EMC SmartFabric OS10 User Guide* available in Dell EMC Networking OS10 Info Hub.

Enable SFS

This information describes how to enable SmartFabric Services. To enable SFS on a switch from the SmartFabric OS10 command-line interface (CLI), use `smartfabric l3fabric enable` command and set a role. In SmartFabric mode, the two leaf or ToR switches are automatically configured as a VLT pair, and the VLT interconnect link (ICL) ports must be physically connected before enabling SFS.

Once you enable SFS on switches and set a role, the network operating system prompts for configuration to reload, then boots in SFS Fabric mode. To apply the changes, enter Yes to confirm and the switch reloads in Fabric mode. The switch is then placed in Fabric mode, and the CLI is restricted to global switch management features and monitoring. SFS Master controls all network configuration for interfaces and switching or routing functions.

Use these SmartFabric OS10 CLI commands to build a leaf and spine fabric:

- On leaf switches:

  ```
  Leaf1(config)# smartfabric l3fabric enable role LEAF vlti icl_ports
  ```

  Example:

  ```
  Leaf1(config)# smartfabric l3fabric enable role LEAF vlti ethernet 1/1/1-1/1/4
  ```

- On spine switches

  ```
  Spine1(config)# smartfabric l3fabric enable role SPINE
  ```

For complete information about how to use SFS commands, see *SmartFabric commands* in the *Dell EMC SmartFabric OS10 User Guide* available in Dell EMC Networking OS10 Info Hub.

SFS User Interface

You can also enable SFS using the SFS User Interface (UI). OS10 switches support SFS UI to set up initial SFS configuration in SFS leaf and spine deployment. The SFS UI is focused on day zero deployment operations and management of the switches in a Layer 3 SFS fabric. For more information about the SFS and SFS UI, see SmartFabric Services in the *Dell EMC SmartFabric OS10 User Guide*.

Fabric creation

This information describes switch discovery, SFS Master, Master advertisement, SFS REST services, Master high availability, preferred Master, SmartFabric, rack or VLT fabrics, default fabric settings, reserved VLANs, default client management network, default client control traffic network, and spanning-tree protocol.

Switch discovery

When SFS is enabled on PowerSwitches, the switches boot in SmartFabric mode, then start discovering each other using LLDP. All discovered switches become part of a single SFS domain, to form a single network domain.

NOTE: For L3 fabric profile, the SFS `Domain ID` is automatically set to 100 and is not configurable in the current release. All directly connected switches join one single domain.

The port where another leaf switch is discovered is configured as a VLT interconnect link (VLTi or ICL), and the port where another spine switch is discovered is configured as an interswitch link (ISL). A switch operating as a spine can only have ISL links to other leaf switches.

SFS uses reserved VLAN 4000 internally to establish communication between switches in a single network fabric. VLAN 4000 is automatically added to all ICL and ISL ports.

SFS Master

SFS uses Keepalive protocol, running on VLAN 4000, to elect one in the fabric as a Master switch. Only a leaf switch can be elected as a Master.

In a single SFS domain, there is only one Master switch at any given time, and the rest of the leaf switches are designated as the backup. A new Master is elected from the backup switches when the Master fails to provide high-availability to the fabric.

NOTE: Spine switches cannot be elected as a Master node within SFS.

Master advertisement

Once a Master is elected, it initiates all applications to automatically build the network fabric. The Master VIP is advertised using mDNS Avahi services for applications to automatically discover the fabric through inband networks.

SFS REST services

The SFS REST service is started on the Master node. Applications consuming or integrating with SFS use this REST service for fabric operations. Communication is performed with the fabric using the IPv6 VIP assigned to the SFS Master, or using the IPv4 out-of-band Management IP of the Master.

A default REST_USER account is created to authenticate all REST queries. The default password is `admin`, and Dell Technologies recommends changing the password through VxRail Manager or OMNI.

NOTE: OMNI communicates with SmartFabric REST Services through REST_USER account only.

Master high availability

SFS uses an internal distributed data store where all fabric configuration is saved. This data is synchronized with all backup switches ensuring the Master, and the backup switches always have the same view of the fabric. With a Master failover, the switch taking over as the Master uses its internal data store to continue fabric operations.

When the fabric is expanded, the newly added switches receive all fabric policies from the SFS Master, once the switches are added to the domain.

Preferred Master

When a Master is elected for a fabric, the switches that are configured as Preferred Master have a higher priority to become the Master switch. If none of the switches are configured as the Preferred Master, any leaf switch can become the Master.

When the fabric is expanded, newly added switches may come up and form a fabric among themselves, and elect a Master before they are connected to the existing fabric. When the new fabric merges with the existing fabric, SmartFabric elects a new Master switch for the combined fabric. If one of the new leaf switches becomes the master, it may overwrite the configuration in the existing fabric. Ensure that the leaf nodes in the existing fabric are set as the Preferred Master before expanding the fabric to prevent the configuration loss.

NOTE: When an uplink is created from SFS UI, preferred master is automatically set. If the uplink is not created using SFS UI, when an uplink is created from OMNI, OMNI sets Preferred Master flag automatically on all configured leaf switches in the fabric if not configured already.

SmartFabric or SFS domain

SmartFabric or SFS domain is interchangeable terminology, and the fabric consists of all switches directly connected to form a single logical network. The L3 fabric is automatically assigned ID 100 and this ID cannot be changed. The fabric name and description are automatically assigned, but can be changed through the SFS user interface.

Rack or VLT fabrics

When two leaf switches are discovered on specified VLTi ports, a VLT is automatically created between the two switches to form a network fabric called the VLT fabric. This VLT fabric is automatically assigned with a fabric ID, a universally unique identifier (UUID).

In a single rack deployment, VLT fabric shows two leaf switches and network fabric shows all the switches. The rack has a network fabric and the VLT fabric represent the same set of switches. In a multi rack deployment, each rack has a VLT fabric, and all the VLT fabrics and the spine switches together form the network fabric.

Default fabric settings

SFS automatically builds the network fabric using industry-standard Layer 2 and Layer 3 protocols.

Reserved VLANs

To build fabric, SFS reserves VLANs 4000 to 4094 for internal use. You are not allowed to use these VLANs for general use.

- **VLAN 4000 — SFS control VLAN** SFS automatically configures VLAN 4000 on all switches that are discovered in the fabric, and uses it for all fabric operations internally. When a leaf or spine is switch is discovered, the ICL or ISL ports are automatically added as tagged members.
- **VLAN 4001 to 4079 — Leaf and Spine connections** SFS automatically sets up the leaf and spine network configuration using eBGP as the underlay routing protocol. SFS uses the reserved VLAN range (4001 to 4079) with automatic IP addressing to set up the peer connections. When a spine switch is connected to the fabric, an ISL is created between the leaf and spine switch. Each ISL link uses a reserved VLAN and the ISL ports that are configured to be the untagged members of this VLAN. IP addresses from the reserved range are used for this VLAN, and an eBGP session is started on the VLAN IP interface.
- **VLAN 4080 — Global untagged VXLAN VLAN** SFS automatically sets up VXLAN overlay networks with EVPN to extend networks between racks in a multi rack deployment. SmartFabric OS10 requires an untagged VLAN on leaf switches for VXLAN traffic handling when using VLT. VLAN 4080 with automatic IP addresses from the reserved range is used for leaf-to-leaf interconnect (ICL) links.
- **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 sets up iBGP peering between a pair of leaf switches directly connected over ICL links. VLAN 4090 with automatic IP addresses from the reserved range is used for enabling iBGP sessions between the VLT peer switches.
- **VLAN 4094 — VLT control VLAN** SFS automatically creates VLAN 4094 on all leaf switches. VLAN 4094 is used for all VLT control traffic between two VLT peer switches. VLAN 4094 is only added on the VLT interconnect links (ICL ports) on leaf switches.

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 customer deployment, these default networks may be changed in the SFS UI using **Edit Default Fabric Settings**. For complete information about SFS UI, see Dell EMC SmartFabric OS10User Guide available in Dell EMC Networking OS10 Info Hub.

NOTE: OMNI uses a docker instance with the default IP address of 172.16.0.1/16. Change the docker private network configuration if the docker IP address conflicts with an existing customer deployment network, see Configure docker private network.

Default client management network

SFS automatically sets up 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 this VXLAN network.

NOTE: The embedded SFS user interface allows you to change this VLAN to a specified VLAN.

Default client control traffic network

SFS sets up a second overlay network that is called *client control network* for VxRail integrated solutions. When a VxRail node is discovered, it is automatically added as a tagged member of this network. SFS also enables the mDNS Avahi service on this network for master advertisement and fabric discovery by integrated solutions. The SFS Master virtual IP for VXLAN network is advertised. The VIP address is `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. Although you can change the VLAN associated with this, it is not recommended to change it for VxRail integrated solution deployments.

Spanning-tree protocol

SFS uses RPVST+ as the default spanning tree protocol to build leaf and spine switches.

Spanning-tree protocol is disabled for VXLAN networks. SFS automatically creates user networks as VXLAN networks in the fabric. For a Layer 2 uplink from the fabric to the external network, the uplink ports in the fabric are configured as VXLAN access interfaces and spanning-tree BPDUs are not sent to the external network.

SFS support for MSTP on L3 fabric: By default, the STP mode is RPVST+. You can change the mode to MSTP once the fabric is built from SFS UI. When you change the mode, the whole fabric goes through a reboot cycle and the new mode is set as MSTP.

NOTE: Changing the mode impacts traffic in the SFS as fabric reboots.

The spanning tree behavior for Layer3 fabric is as follows:

- STP is enabled on Cluster control VLAN (VLAN 4000). The spine switches are configured to take over the STP root role.
- STP is disabled on all inter leaf-spine VLANs and leaf-leaf VLAN (4001-4091).
- STP is enabled on all user created VLANs.
- STP is disabled on server facing port.

NOTE: VLANs used for setting up the leaf and spine eBGP peering are automatically set up to prevent loops while having nonblocking connections between the leaf and spine switches.

SFS and OMNI supported solutions

OMNI 2.0 with the SmartFabric Services OS10 release supports the following qualified solutions. See the Solutions Support Matrix for the latest supported versions for all the qualified solutions. :

Table 4. Qualified solutions

Qualified Solutions	Dynamic discovery	Onboarding type	vCenter/Day 2 automation
VxRail	Yes	Automatic	Yes
PowerEdge MX	NA	NA	Yes
PowerStore X (ESXi)	Yes	Import from Fabric or vCenter	Yes
PowerStore T	Yes	Import from Fabric	No
Isilon front-end/PowerScale	No	Manual	No
Other devices running ESXi	No	Import from vCenter or Manual	Yes
Other devices running Windows or Linux-based Operating Systems	No	Import from Fabric	Yes

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

NOTE: Other devices can be supported provided they meet the industry Ethernet standards and are compatible with SmartFabric-enabled switches.

Dynamic Discovery - Devices that support dynamic discovery send a Dell-specific LLDP TLV. Supported devices are automatically populated in the SFS UI and OMNI by MAC address, switch, and switch port number for onboarding to the fabric. Devices that do not send the Dell-specific LLDP TLV must be manually added to the fabric.

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

vCenter/Day 2 Automation - Port groups that are created in vCenter are automatically applied to the applicable host-connected ports on the switch. The host must be running ESXi, added to the vCenter, and have a server profile that is created in OMNI. For the automation to work, register OMNI with the vCenter and ensure to start the respective OMNI vCenter automation services.

SFS personalities

Two types of SFS:

1. SFS for leaf and spine—supported on selected S-series and Z-series PowerSwitches.
2. SFS for PowerEdge MX—supported on PowerEdge MX switches.

SFS for leaf and spine has two personalities:

- **SFS VxRail L2 single rack personality**—enables an automated single rack network fabric (L2 fabric profile) for VxRail clusters. Use the L2 personality for the existing fabric deployments. For more information about configuring VxRail L2 single rack personality, see *VMware Integration for VxRail Fabric Automation SmartFabric User Guide, Release 1.1, September 2019*. For new SmartFabric deployments, it is recommended to use the L3 leaf and spine fabric personality for future expansion.
- **SFS L3 leaf and spine fabric personality**—enables a multi rack data center network fabric offering flexibility to start with a L3 single rack (L3 fabric profile), and expand to a multi rack solution on demand. The L3 personality is integrated with VxRail to enable single-site, multi rack VxRail deployments allowing VxRail nodes to be easily deployed in any rack without complex underlay network configuration.

OpenManage Network Integration (OMNI) enables fabric management and zero-touch automation for:

- SFS L3 leaf and spine fabric personality
- SFS VxRail L2 single rack personality

Table 5. SFS personality comparison

SFS VxRail L2 single rack personality	SFS L3 leaf and spine fabric personality
Network fabric with two ToR switches in a single rack cannot be expanded beyond a single rack.	Network fabric with up to 20 switches in a leaf and spine design that can start with a single rack, and extend up to eight racks. If you want to deploy a L3 single rack fabric, enable only leaf switches in the rack without spine. Add spine to the L3 single rack to form a L3 multi rack leaf and spine fabric.
All VxRail SmartFabric deployments earlier to SmartFabric OS10 10.5.0.5 releases.	All new SmartFabric deployments with SmartFabric OS10 10.5.0.5 or later.
Enabled through shell commands with fixed parameters.	Enabled through standard SmartFabric OS10 CLI commands with role and VLTi ports for leaf as fixed parameters. Enable SFS using SmartFabric UI also. For more information about SFS UI, see *Dell EMC SmartFabric OS10 User Guide*.
Default uplink and jump box port that is created as part of SmartFabric initialization, and cannot be modified after enabling SFS as part of Day 2 operations.	The network fabric is created as part of SmartFabric initialization. Uplinks and jump box port must be created through the embedded SFS user interface or OMNI and are fully customizable as part of Day 2 operations.
In VxRail deployment and Day 2 operations, all networks that are created during initialization are VLAN backed network with customer router acting as the gateway.	Networks that are created during initialization and as part of VxRail deployment and vCenter integration are VXLAN

Table 5. SFS personality comparison (continued)

SFS VxRail L2 single rack personality	SFS L3 leaf and spine fabric personality
	stretched networks for single rack deployments. VLAN-based networks in a rack can be created through OMNI.
Existing deployments when upgraded to SmartFabric OS10 10.5.0.5 continue to run in L2 mode. L3 fabric capabilities are not available.	Migration from VxRail L2 personality to L3 fabric personality is not available with SmartFabric OS10 10.5.0.5 version, and will be available in a future release.

NOTE: Dell Technologies recommends enabling all new deployments with L3 leaf and spine fabric personality. VxRail SmartFabric deployments using older VxRail L2 single rack personality cannot be upgraded to the new L3 leaf and spine fabric personality automatically. A migration workflow will be available in a future release to allow existing deployments to expand to a multi rack solution.

4

OpenManage Network Integration

OpenManage Network Integration (OMNI) enables configuration of SmartFabric Services (SFS) that integrates with VMware vCenter for fabric automation of the physical network infrastructure corresponding to the virtual network operations within vCenter. OMNI also serves as a front-end management application for managing one or more service instances, enabling administrators to manage and operate one or more network fabrics that are deployed with SFS.

OMNI virtual appliance

The OMNI virtual appliance is delivered as an open virtual appliance (.ova extension) file. Deploying an OMNI OVA template allows you to add preconfigured OMNI virtual machines to vCenter Server or ESXi inventory.

The OMNI OVA file can be downloaded from the Dell EMC OMNI for VMware vCenter support portal. OMNI virtual machine deployment is tested and supported only on the VMware ESXi hypervisor, even though it is expected that the OVA could be deployed in other x86 hypervisors.

OMNI deployment

Deploying an OVA template is similar to deploying a virtual machine from a template. You can deploy an OVA template from any local file system accessible from the vSphere web client, or from a remote web server.

Table 6. OMNI deployment

OMNI VM system requirements	**vCenter Server Network (OMNI VM Network 1 - ens160)**	**VxRail Management Network (OMNI VM Network 2 - ens192)** ***Optional in non-VxRail deployment***	**OMNI access**
• Virtual hardware version: vmx-14 • Compatible: ESXi 6.7 and later • 4 virtual CPUs; 8 GB memory; 80 GB hard disk	Out-of-band (OOB) management network • Provides reachability to DNS, default gateway, and where OMNI obtains the IP address or hostname • Provides reachability to Management network (vCenter IP address or FQDN, SmartFabric Management IP address or hostname)	In-band link-local network—Provides reachability to SmartFabric link-local network for IPv6 VIP reachability	• vCenter HTML5 (/ui) plug-in; click OpenManage Network Integration link. • OMNI stand-alone UI: https://*OMNI_IP or FQDN*/ using `admin` user • SSH to OMNI VM IP address or FQDN as `admin` user • OMNI VM console using vCenter or ESXi `admin` or `root` user
	VxRail default: vCenter Server network	VxRail default: VxRail Management network	

NOTE: Even when OMNI is deployed in-band, Dell Technologies recommends setting up connectivity with the out-of-band Management network of the switches in the network fabric to separate management traffic with user data traffic, and also to enable faster image downloads to the switches.

Maximum supported instances

A single OMNI VM instance supports:

Table 7. Number of supported instances

Entities	Number of instances supported by OMNI
vCenter	8
SmartFabric instances	15
OME-Modular instances	1
NSX-T Manager	1

Install OMNI virtual appliance using vCenter

This information describes how to deploy the OMNI appliance on a VMware ESXi hypervisor using the OMNI OVA file, and create a virtual machine (VM).

NOTE: The OMNI plug-in or SmartFabric Services user interface does not provide localization.

When upgrading from older version to 2.0, follow the instructions that are provided in Upgrade OMNI appliance.

Download and install OVA

1. Download the OMNI release package from OpenManage Network Integration support locally, and extract the OVA image and README files from the release package .
2. Validate the code signed OVA image according to the instructions in README file in the release package. If the signature is invalid, contact Dell EMC Technical Support for a valid signed image.
3. In the vSphere Client, select **Hosts and Clusters**, right-click the cluster that the plug-in must manage, and select **Deploy OVF Template**.

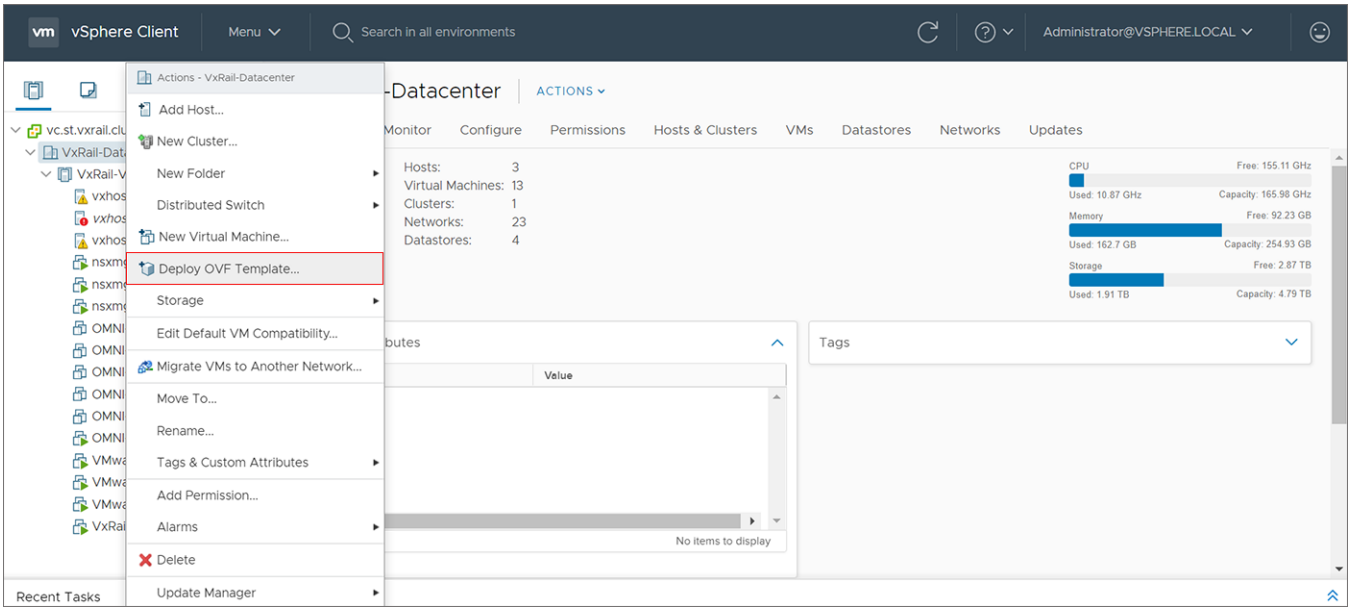

4. Select **Local file**, click **Choose Files**, select the OMNI ova file from a local source, and click **Next**.

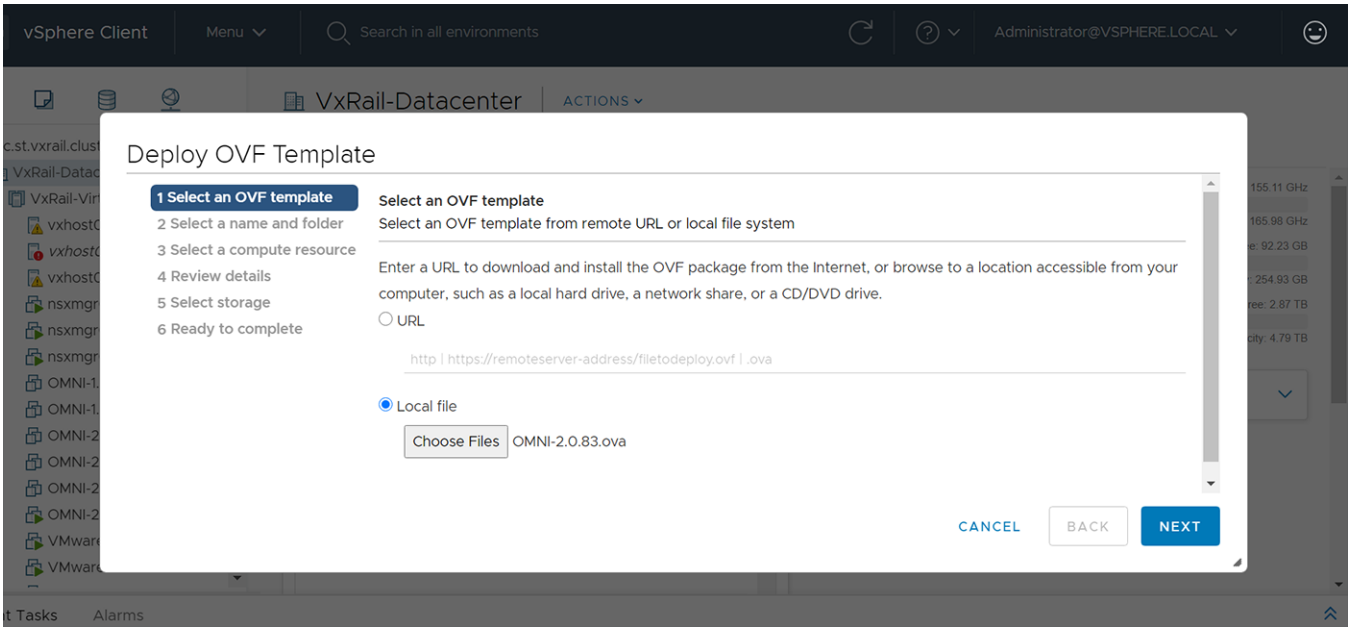

5. Select a name and folder for the VM, and click **Next**.

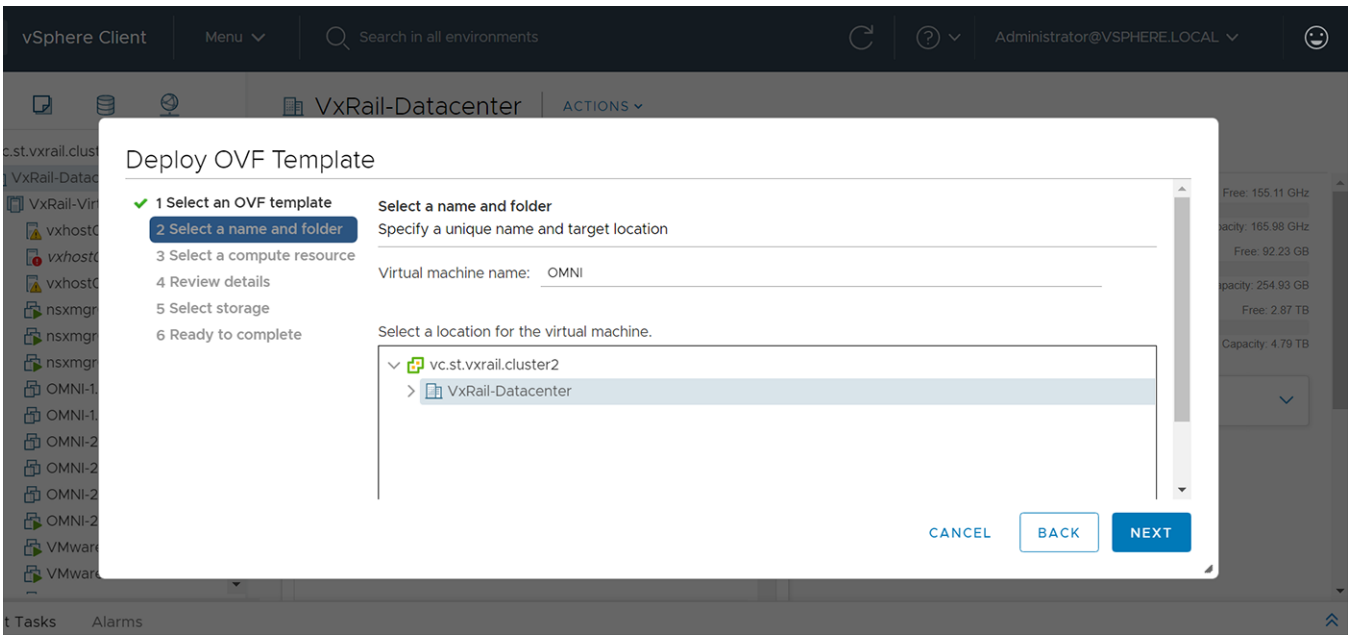

6. Select the destination compute resource, and click **Next**.

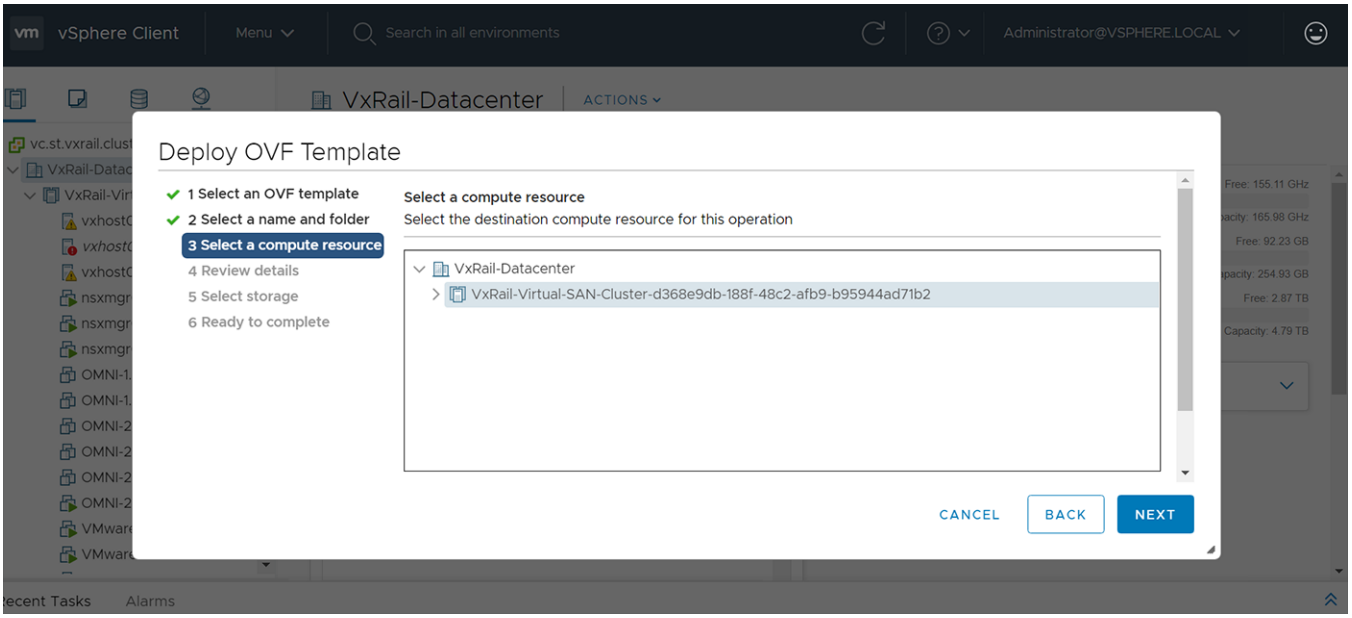

7. Review and verify the template details, and click **Next**.

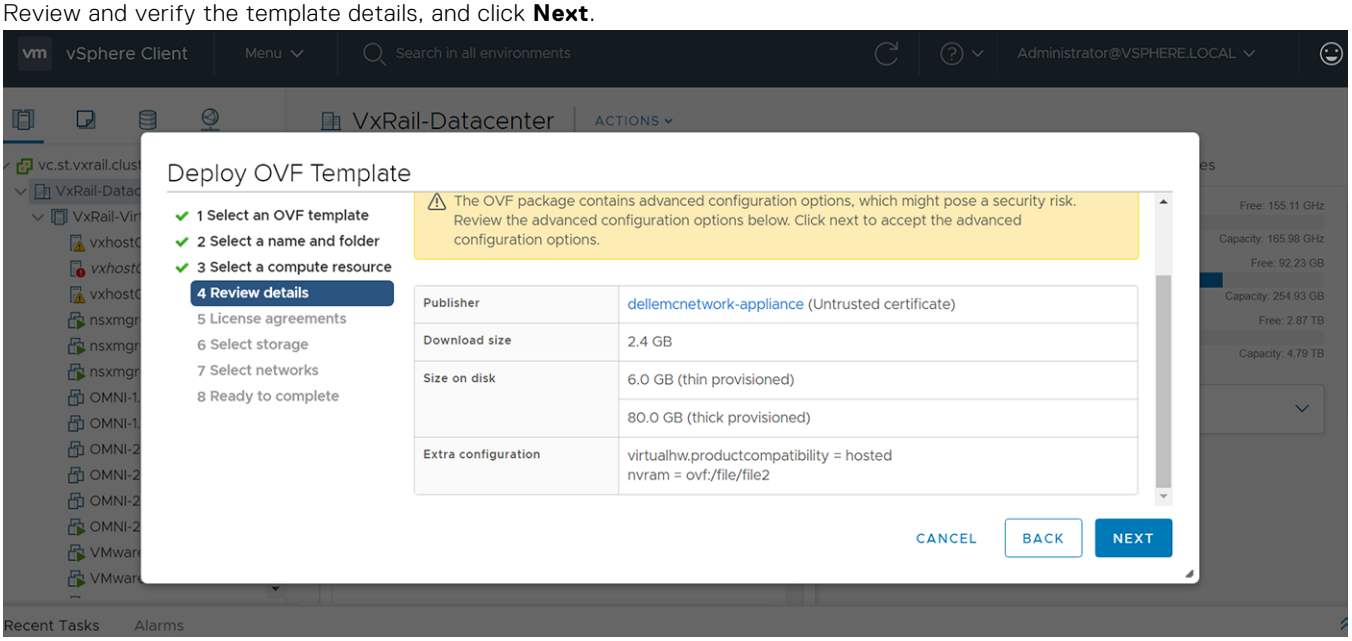

8. Accept the end-user license agreement (EULA), and click **Next**.

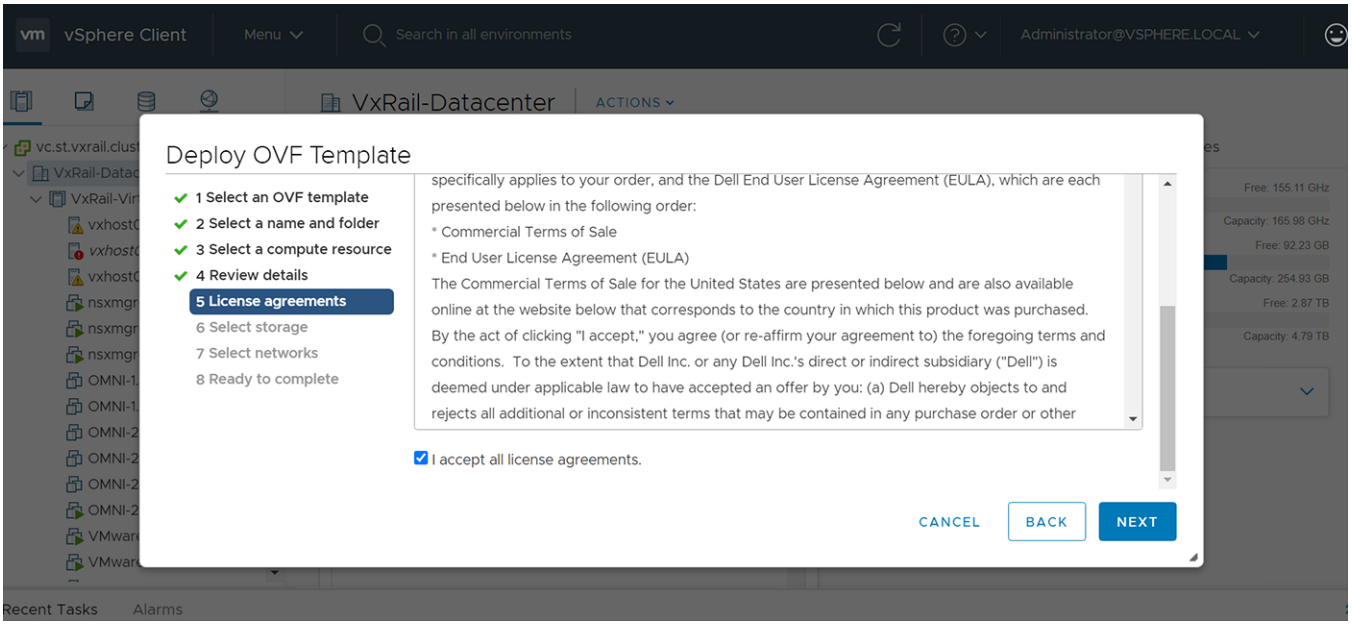

9. Select the VSAN datastore for the configuration and disk files, and click **Next**.

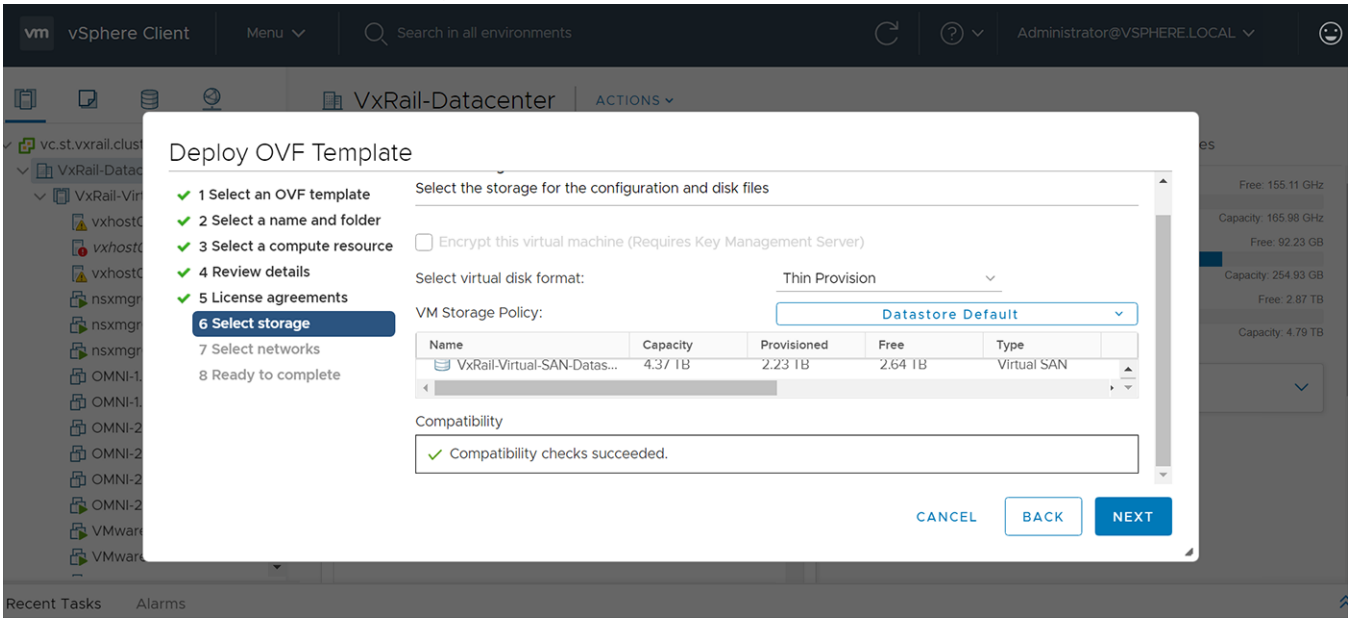

10. Select a destination network for each network source, and click **Next**. The VxRail Management Network must be assigned to the **VxRail internal Management network**. The default VLAN ID for this network is **3939**. The vCenter Server network must be connected to the port group where the vCenter Server is reachable for deployment of

the OMNI plug-in. **If you are using a standalone generic ESXi host deployment, you can skip this step.**

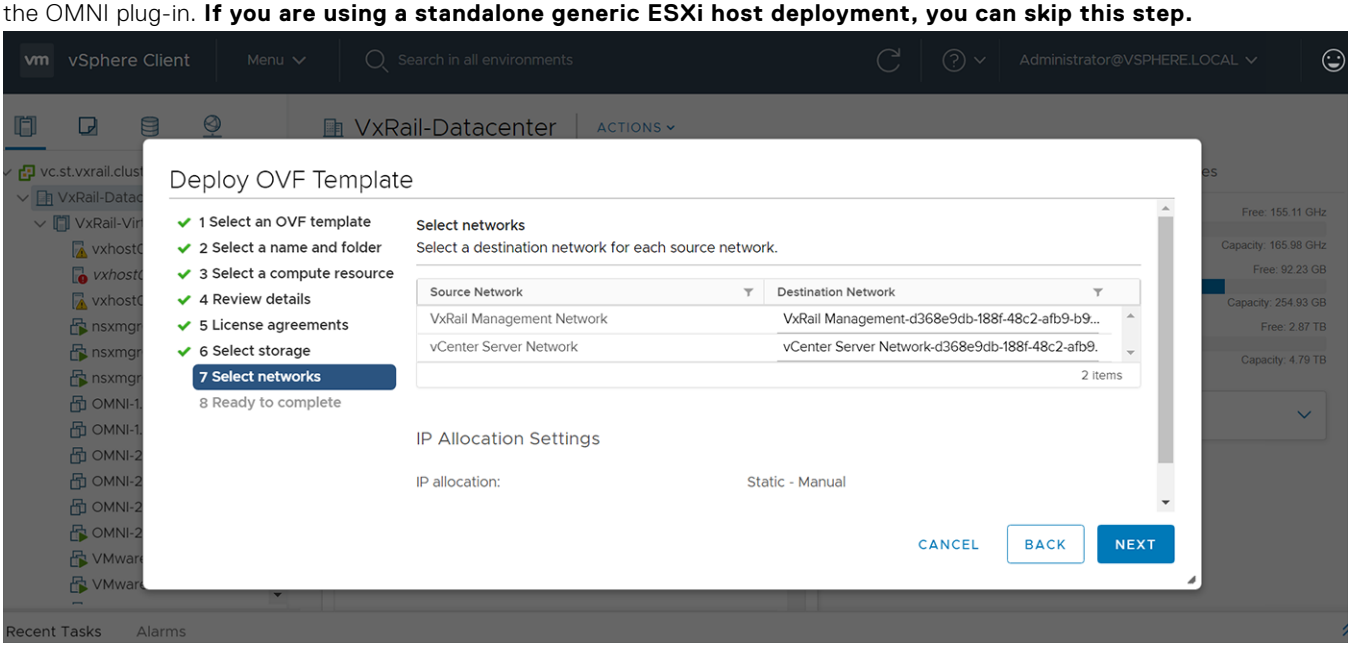

NOTE: Ensure that the source and destination networks are mapped properly. Any misconfiguration may cause connectivity issue between vCenter and OMNI.

11. Click **Finish** to start creation of the VM.

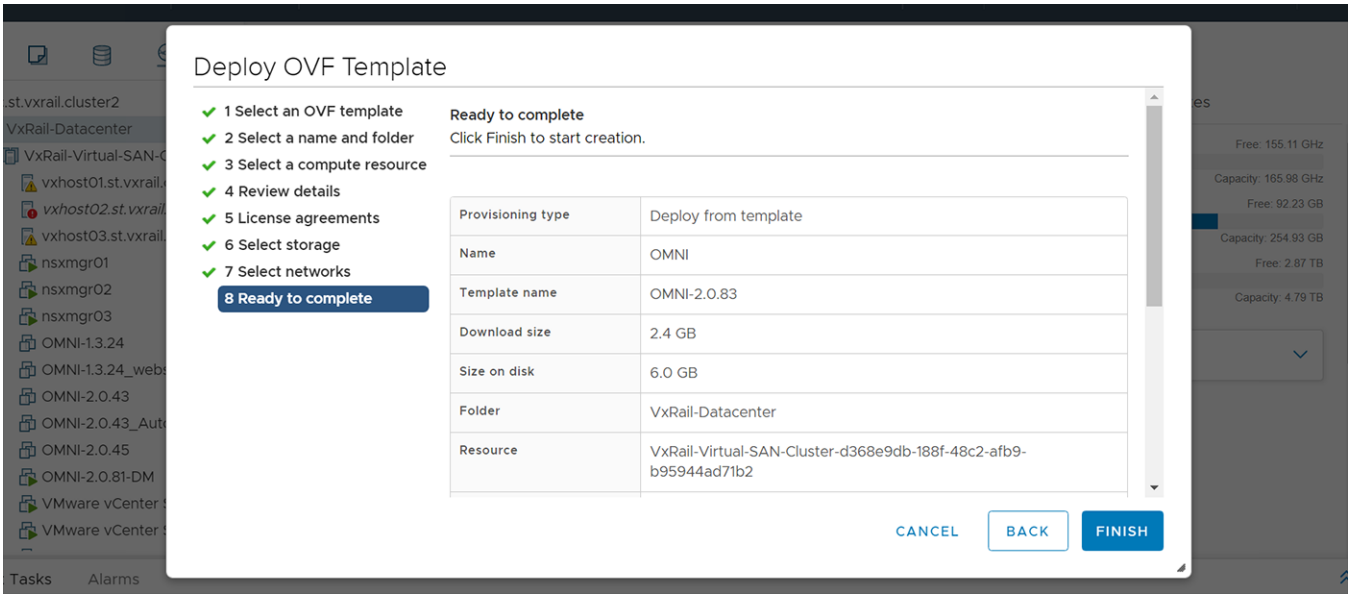

Power on OMNI VM

1. Click **Recent Tasks** and scroll to the bottom of the window to view the status, and wait for the deployment to finish.

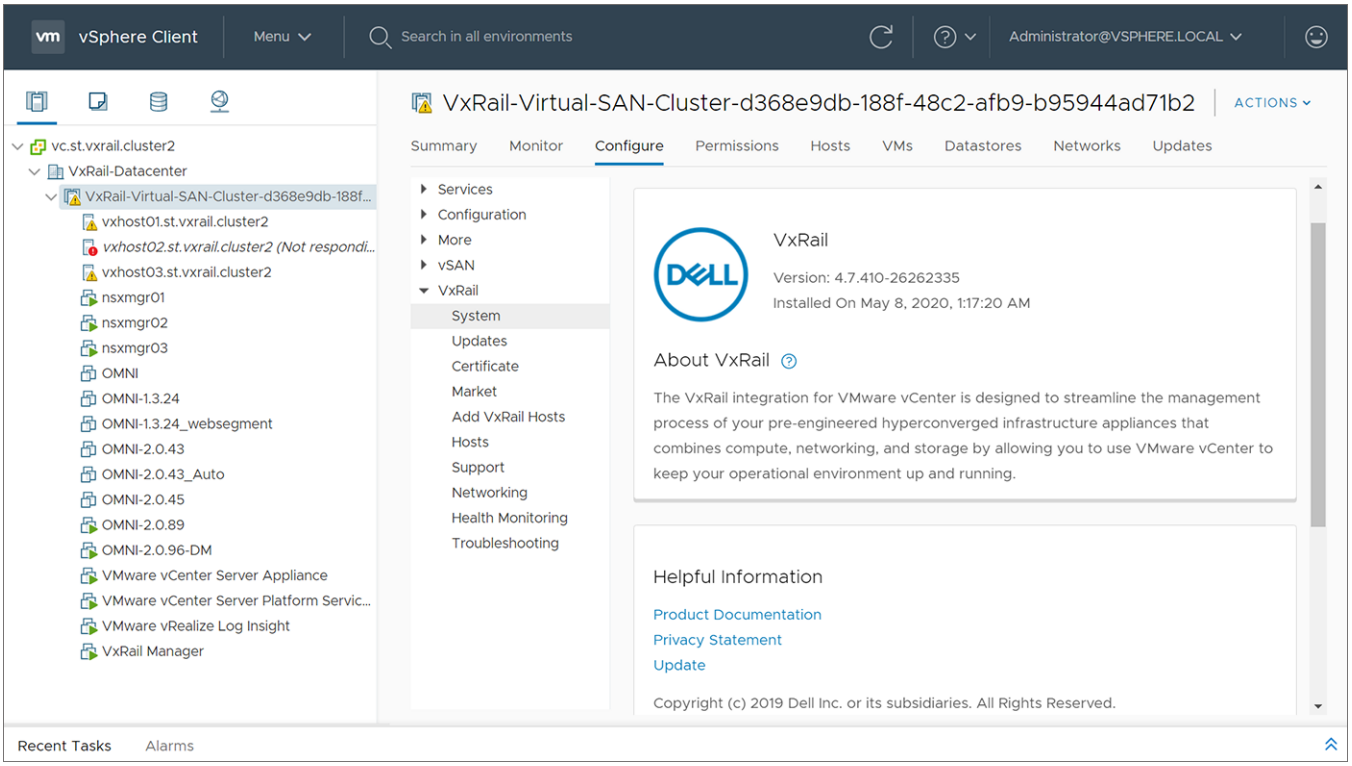

2. Select the OMNI VM you want to power on, and select **Actions** > **Power** > **Power On**.

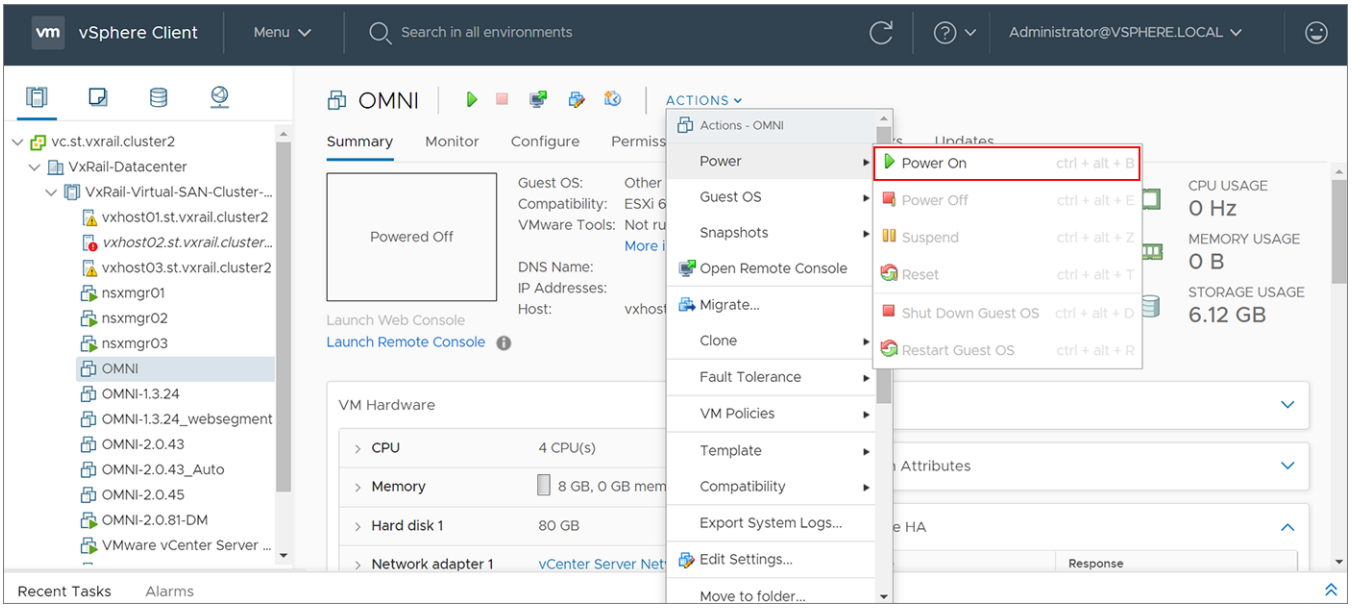

3. Select **Launch Web Console**.

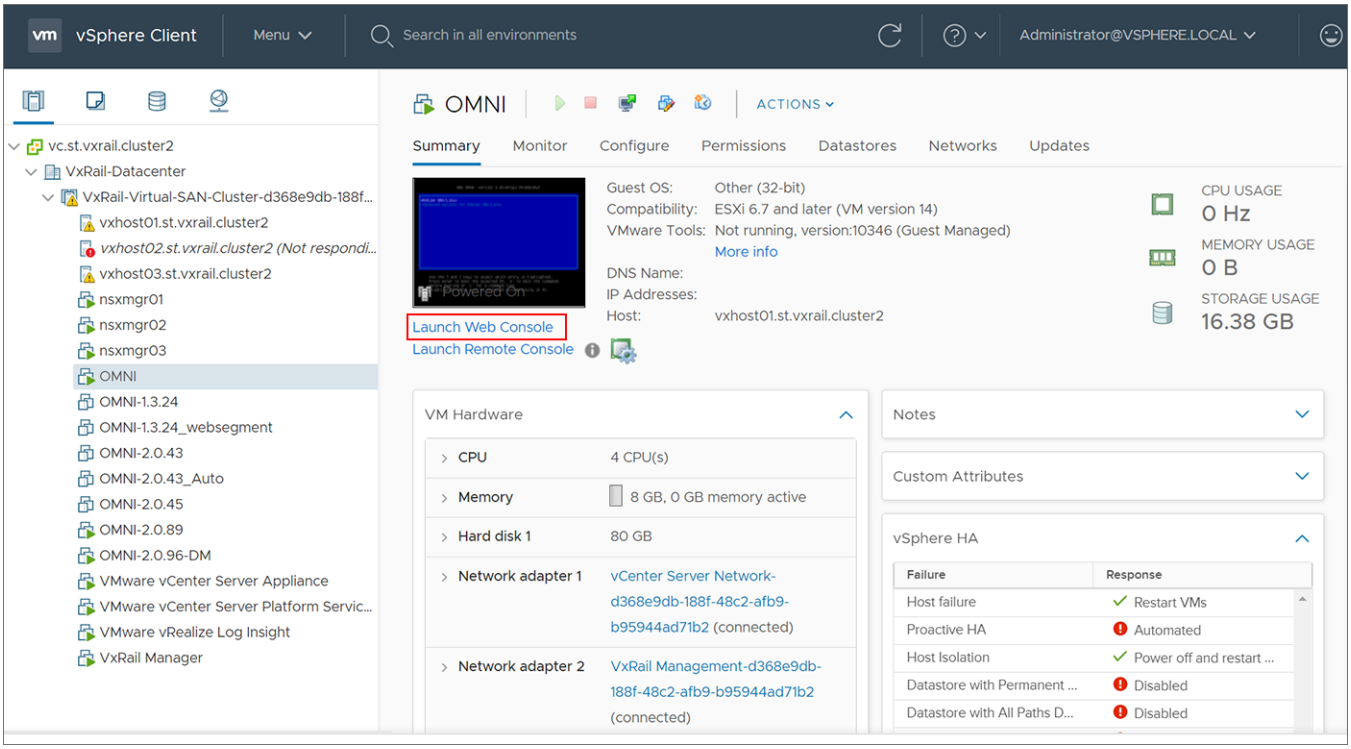

Upgrade OMNI appliance

This section explains how to upgrade the OMNI appliance in two ways.

When upgrading OMNI VM from 1.3 to 2.0 version, you can install the OMNI .ova file using new installation or upgrade OMNI.

Upgrade OMNI - new installation

Follow the below steps when upgrading OMNI appliance from older version (1.1 or 1.2) to 1.3 and later versions.

1. **Prerequisite**

   Save the following details:
   - IP address or hostname of the SmartFabric instances that are manually added in the OMNI VM.
   - IP address or FQDN information of all the vCenters that are registered with the OMNI VM.
   - IP address or hostname of the OMNI VM.
   - Details of the `ens192` and `ens160` interface settings.
2. Unregister the vCenter from OMNI VM, see here.
3. Shut down the older OMNI VM.
4. Deploy the new OMNI VM using the latest OMNI OVA file, see Create OMNI virtual appliance.
5. Configure the OMNI VM with the documented settings and complete the full setup, see Set up OMNI.

Upgrade OMNI

You must be in the OMNI VM console to use these steps and only applies to the OMNI minor release upgrade. You can also upgrade OMNI from 1.3 to 2.0 version using this upgrade workflow.

To upgrade OMNI appliance:

1. Download the OMNI upgrade image from the Dell EMC Support portal and store the image on an SCP server. Check the existing version.
2. From the OMNI VM console, select the option **5. Upgrade Appliance**.

```
################################################################################
    Welcome to Dell EMC OpenManage Network Integration (OMNI) management
################################################################################

    Menu
    ------------------------------------------------
0.  Full setup
1.  Show version
2.  Interface configuration menu
3.  OMNI management service menu
4.  Password/SSL configuration menu
5.  Upgrade appliance
6.  Reboot appliance
7.  Show EULA
8.  Logout

Enter selection [0 - 8]: 5_
```

The display lists all the applications which can be upgraded along with the old and new versions. Upgrading requires restarting the services.

3. Enter the SCP server IP address or hostname, username, and the path to the upgrade .zip file and password.

```
################################################################################
    Welcome to Dell EMC OpenManage Network Integration (OMNI) management
################################################################################

    Menu
    ------------------------------------------------
0.  Full setup
1.  Show version
2.  Interface configuration menu
3.  OMNI management service menu
4.  Password/SSL configuration menu
5.  Upgrade appliance
6.  Reboot appliance
7.  Show EULA
8.  Logout

Enter selection [0 - 8]: 5
2020-11-19 05:20:52 INFO  [setup.sh] Getting the upgrade file
Remote SCP server IP/hostname: 100.104.26.58
Username: admin
Path to the upgrade zip file: /tmp/OMNI-upgrade-2.0.83.zip_
```

4. Verify all information, then enter **Y** to continue.

5. Verify the OMNI version.

```
################################################################################
    Welcome to Dell EMC OpenManage Network Integration (OMNI) management
################################################################################

    Menu
    ------------------------------------------------
0.  Full setup
1.  Show version
2.  Interface configuration menu
3.  OMNI management service menu
4.  Password/SSL configuration menu
5.  Upgrade appliance
6.  Reboot appliance
7.  Show EULA
8.  Logout

Enter selection [0 - 8]: 1
OMNI appliance version ..........(2.0.68)
OMNI vSphere client plugin version .............(2.0.83)
press [enter] to continue...
_
```

6. After you upgrade OMNI, close the active OMNI browser. Open a new browser and log in to OMNI to see the new or upgraded UI changes.
7. Unregister and register the vCenter again using OMNI UI, see Register vCenter with OMNI.

Before upgrading to OMNI from 1.3 to 2.0 version using this upgrade workflow, change the hardware profiles for the VM. To change the hardware details, follow these steps:

1. Go to vCenter and shut down the OMNI VM.

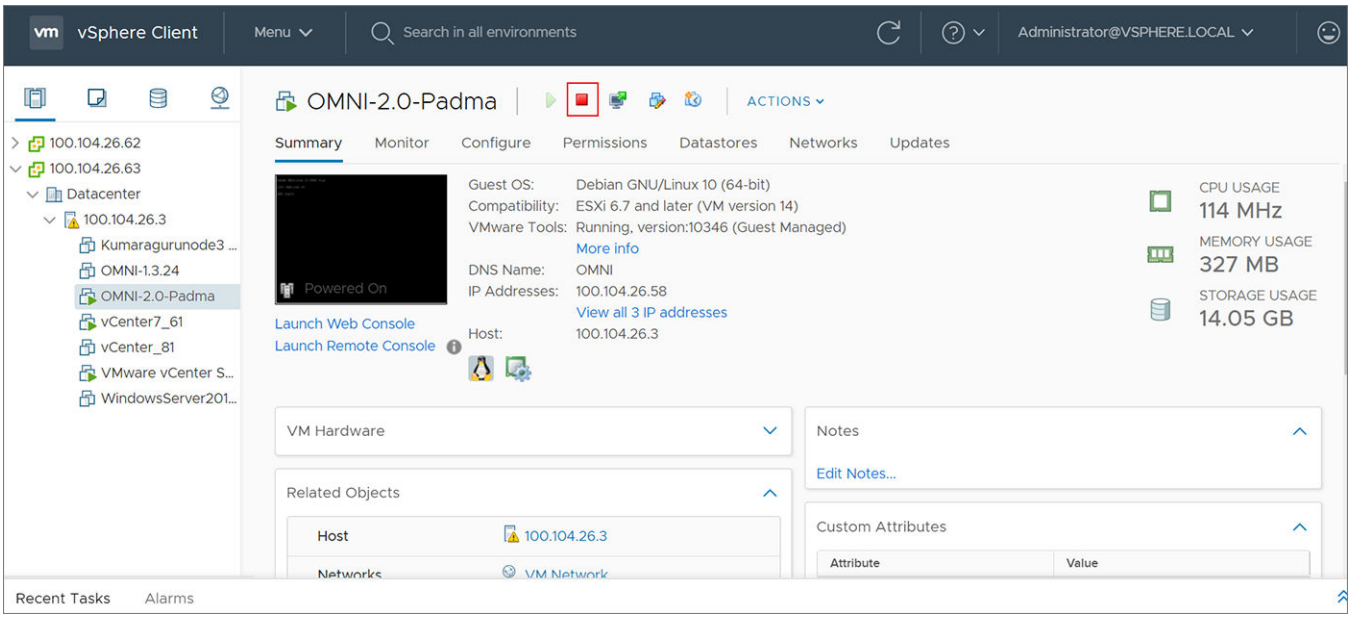

2. Click the **Edit Settings** menu from the OMNI VM page.

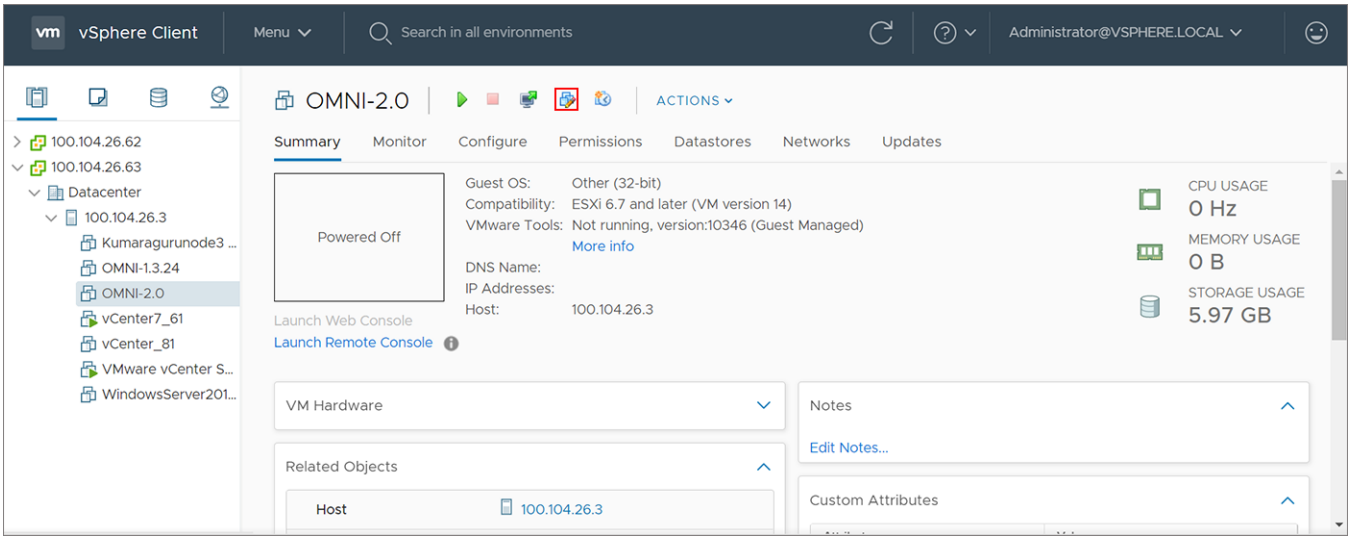

3. Change the **Memory** and **Hard disk 1** settings. Set Memory to 8 GB and Hard disk to 80 GB.

NOTE: When you upgrade OMNI VM using .ova file, you do not have to change these settings as it is installed automatically with the above configurations.

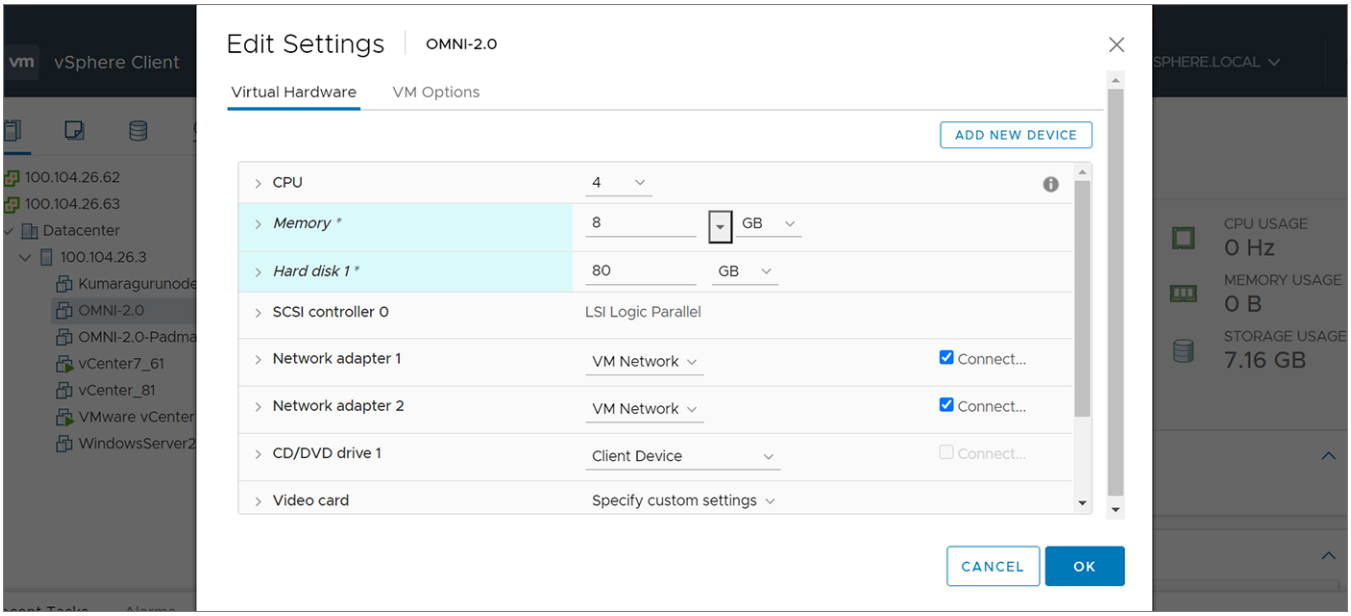

4. Power On the OMNI VM after setting the required configurations.

Set up OMNI

This information describes how to log in to the VM console, and also explains the OMNI vCenter setup.

Log in to VM console

Configure OMNI through the VM console after completing the authentication step. By default, the VM console automatically closes after 10 minutes, but can be customized.

1. Enter `admin` for both the default username and password.

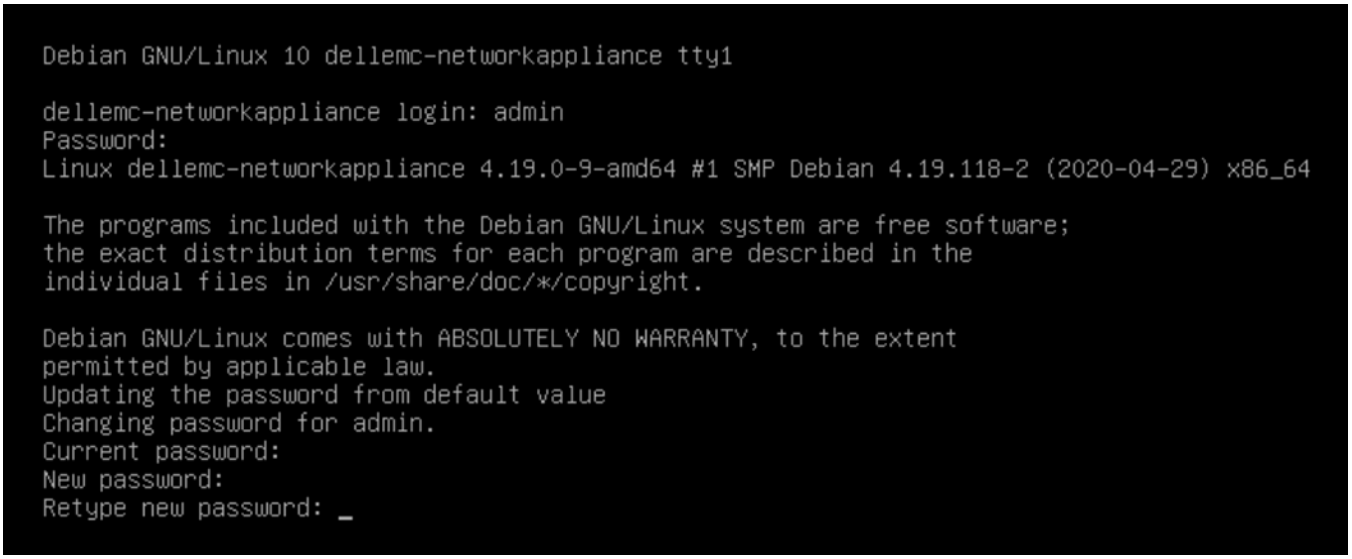

2. If it is a first-time login, the system prompts for password change.

   After the passwords are successfully updated, self-signed certificates are created. You can change the certificates later with OMNI management menu options, see Generate and Install SSL certificate.

NOTE: The sudo password is the same as the password set for the `admin` user.

NOTE: Root user is disabled by default. To set the password to enable **root** user, use the OMNI VM console CLI menu. You can only access root user through the console.

Setup OMNI

This information describes how to set up the OMNI appliance with the required network interface configurations.

NOTE: The OMNI initial configuration setup can be performed using the vCenter OMNI VM console only.

Dell Technologies recommends checking the docker private network setting before setting up OMNI to avoid any conflict with any of the external networks to which OMNI is connected.

NOTE: OMNI default docker private subnet is 172.16.0.1/16. Dell Technologies recommends using 172.16.0.1/24 (x.x.x.x/24) IP address for docker private network setting.

The conflicts occur when:

- The ens160 and ens192 interfaces have IP addresses assigned from the docker private network (172.16.0.0/16).
- Any external entity such as vCenter instance, SFS instance, OME-Modular, NSX-T has IP address that is assigned in the docker private network.
- OMNI is connected to a larger network in which one or more subnetworks IP range overlaps with the docker private network.

If there is a network conflict, OMNI cannot communicate with the other entities. To avoid the conflict, change the docker private network setting, see Configure docker private setting.

Network interface profile configuration

1. Select **0. Full Setup**.

```
################################################################################
    Welcome to Dell EMC OpenManage Network Integration (OMNI) management
################################################################################

    Menu
    ------------------------------------------------
0.  Full setup
1.  Show version
2.  Interface configuration menu
3.  OMNI management service menu
4.  Password/SSL configuration menu
5.  Upgrade appliance
6.  Reboot appliance
7.  Show EULA
8.  Logout

Enter selection [0 - 8]:
```

2. Select **Edit a connection**, then click **OK**.

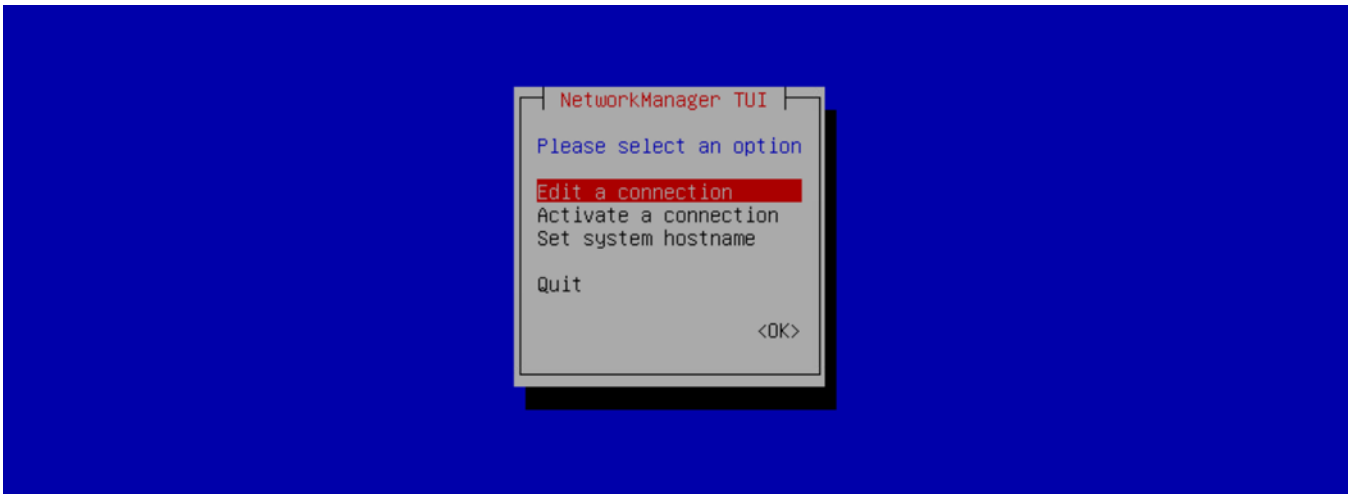

CAUTION: Edit a connection menu displays edit option of Bridge interface `docker0` and `Veth` interfaces, apart from ens160 and ens192. Do not modify any configuration of the `docker0` or `Veth` interfaces as it can lead to OMNI appliance failure or unexpected OMNI behavior.

3. Select **Wired connection 1**, then click **Edit**.

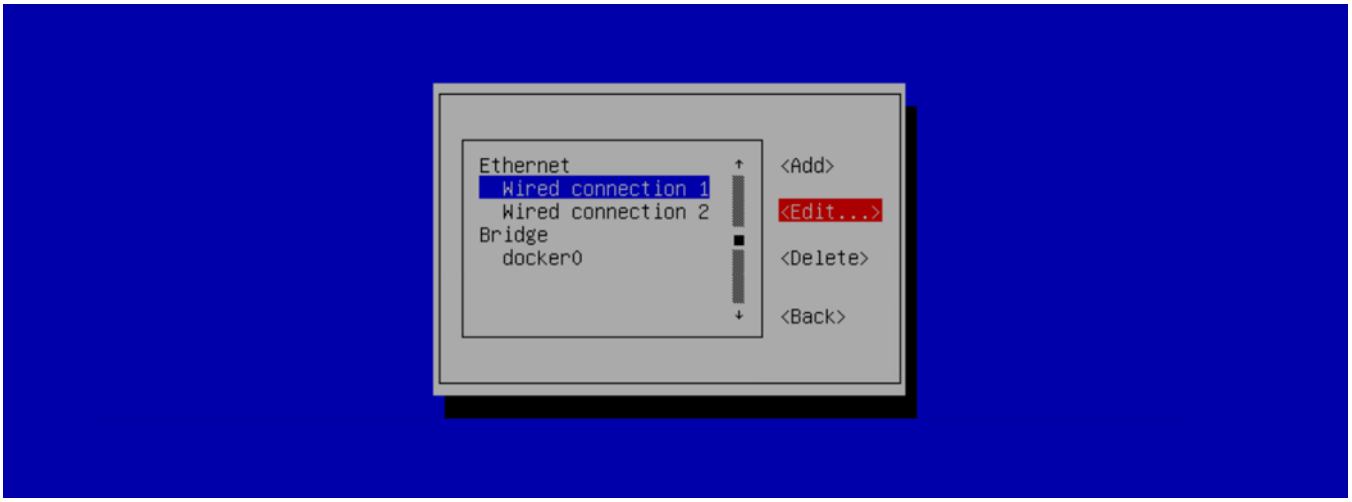

4. Verify Ethernet (ens160) is connected to the vCenter reachable network, then change the Profile name to **vCenter Server Network**.

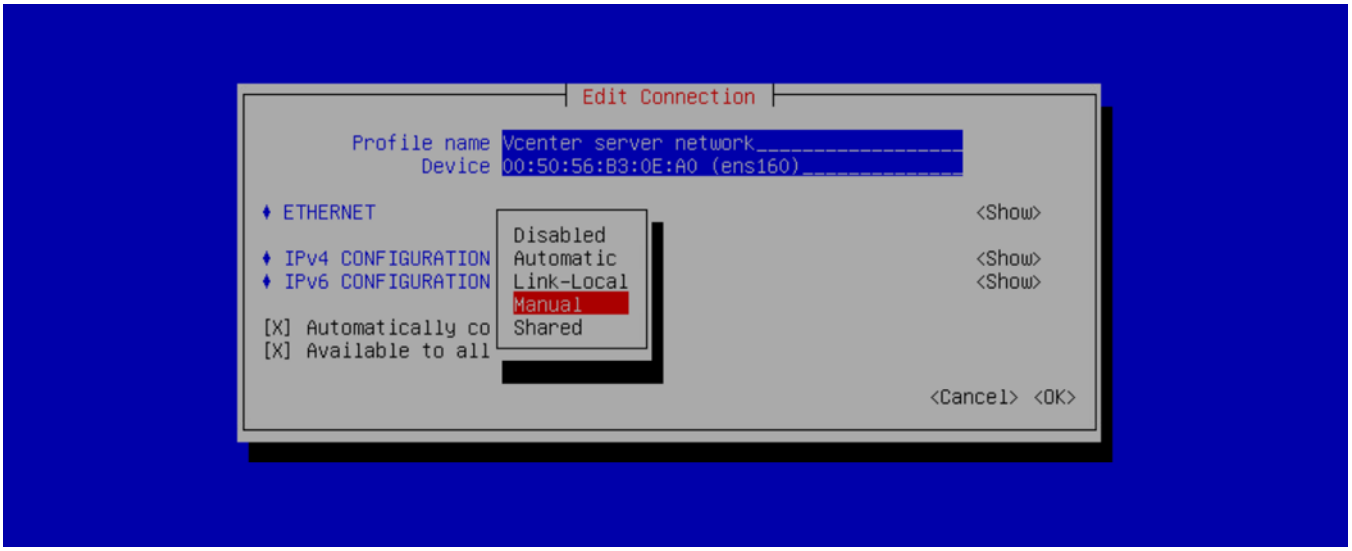

5. Change the IPv4 configuration from Automatic to Manual from the drop-down.

NOTE: If you are using a stand-alone generic ESXi host deployment and if DHCP services are running on the Management network subnet, use the default IPv4 vCenter server network configuration which uses automatic IP address assignment using DHCP. During this scenario, set the IPv4 configuration to Automatic.

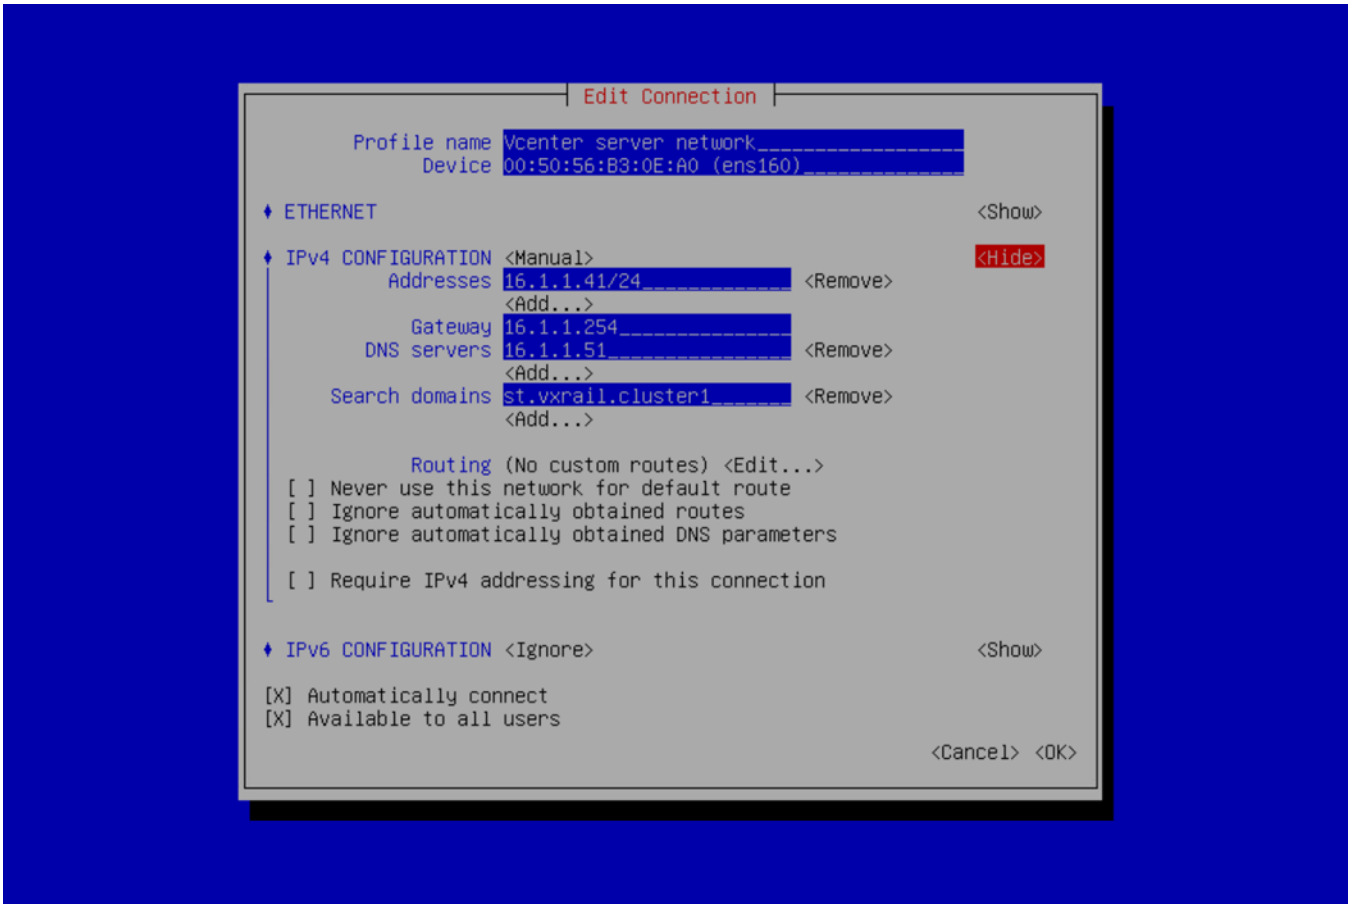

6. Click **Show** to the right of IPv4 configuration, then click **Add**.
7. Set the Manual IPv4 address with subnet mask information, Gateway address, DNS servers, Search Domains, then click **Edit** to the right of Routing.

   NOTE: Ensure that the IPv4 address is set with subnet mask in the prefix-length *(/xx)* format.

8. Select **Ignore** for the IPv6 configuration and click **OK**.

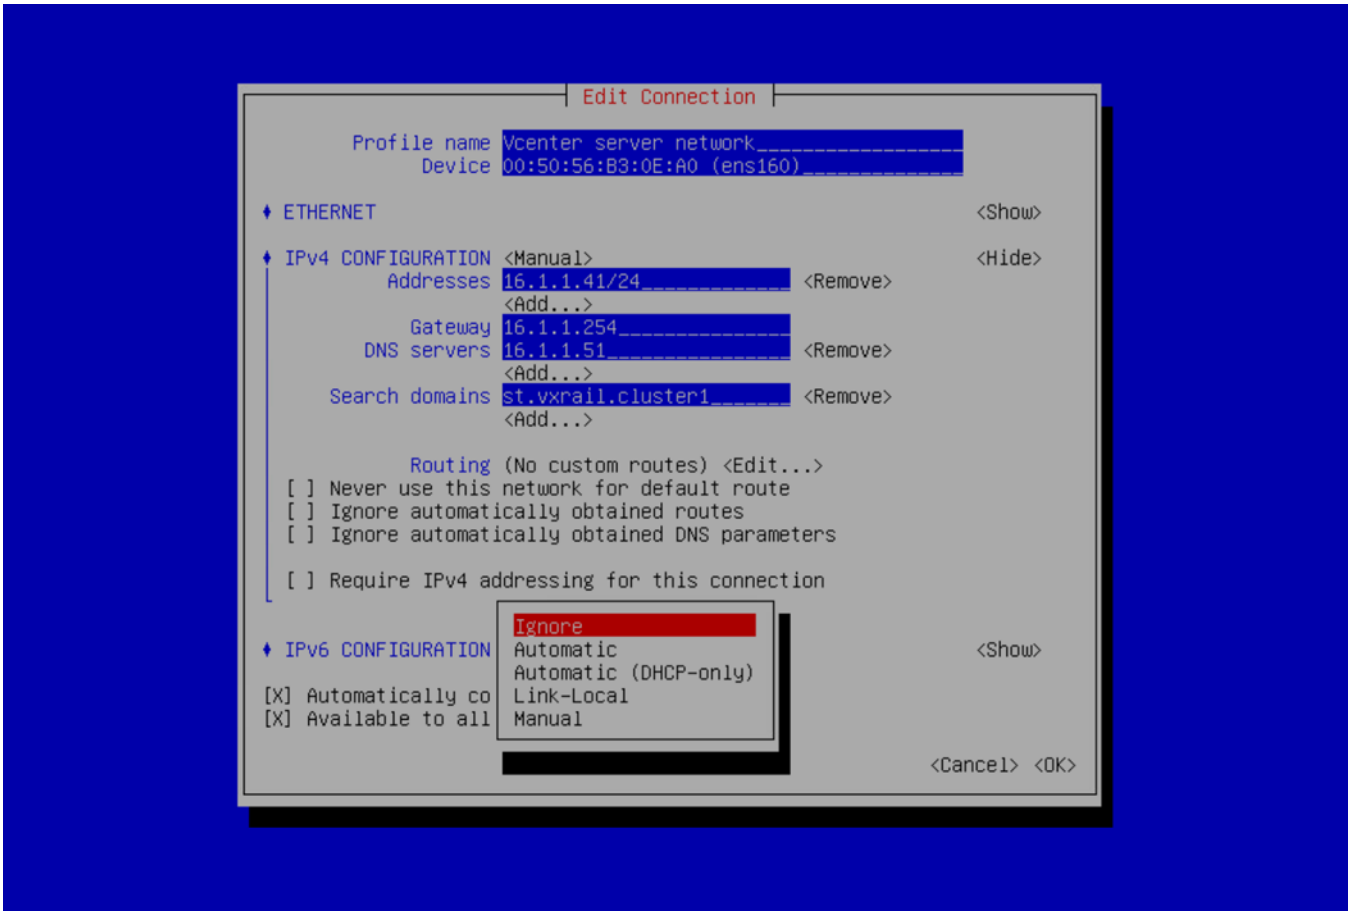

You are now ready to continue configuration.

NOTE: If you are not connecting the OMNI VM to a SmartFabric local-link, ignore this part as it not applicable and you are ready to activate the connection profile.

1. Select **Wired connection 2**, and click **Edit**.

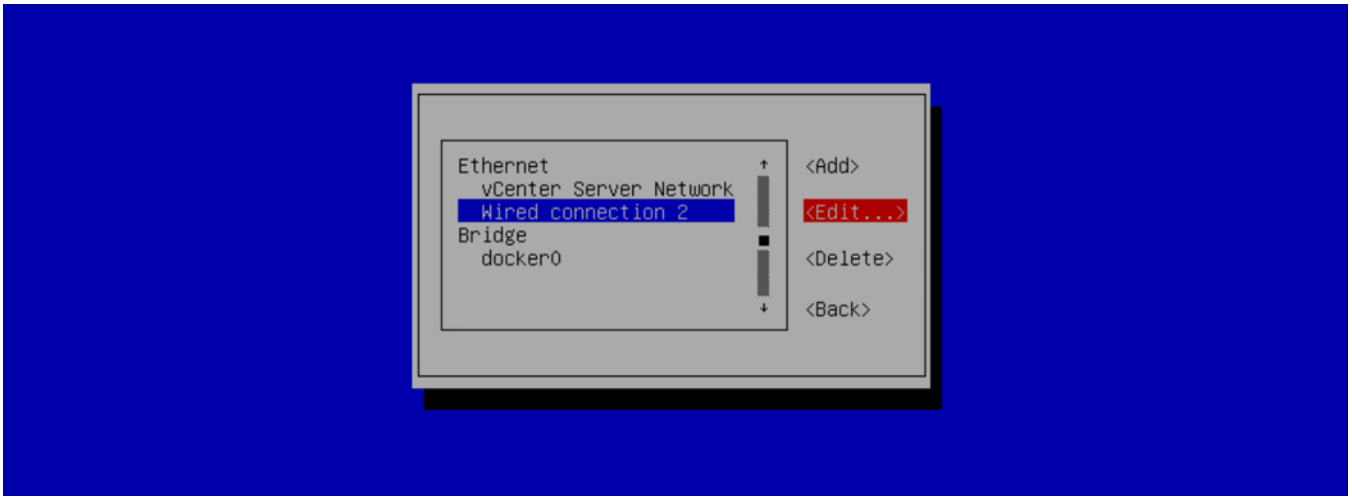

2. Rename Profile name to **VxRail Mgmt Network**.

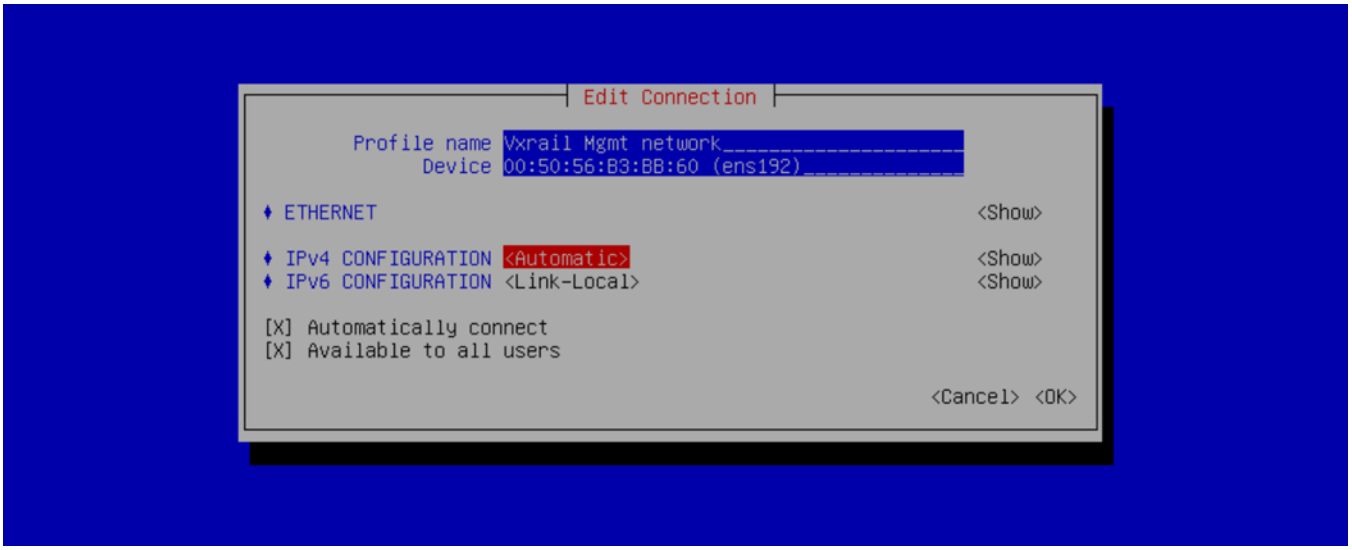

NOTE: The VxRail Mgmt network (ens192) setting is relevant only for VxRail deployment with IPv6 autodiscovered instance. Configuring ens192 interface is not required for non-VxRail environment.

3. Select **Disabled** for the IPv4 configuration.

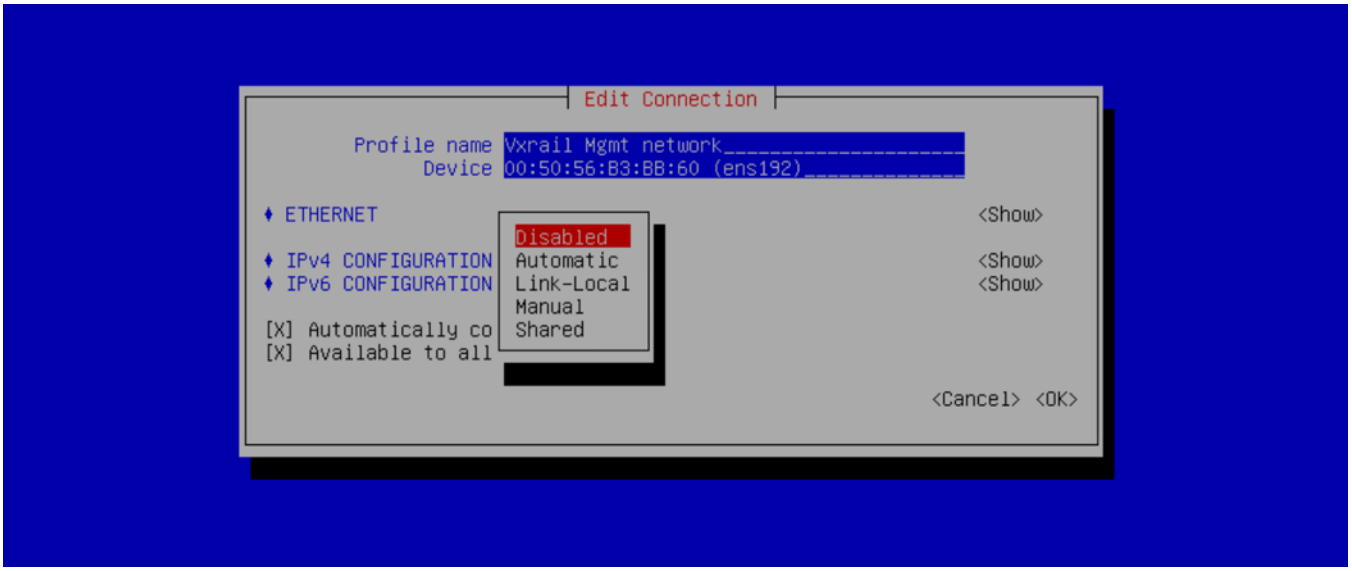

4. Select **Link-Local** for the IPv6 configuration.

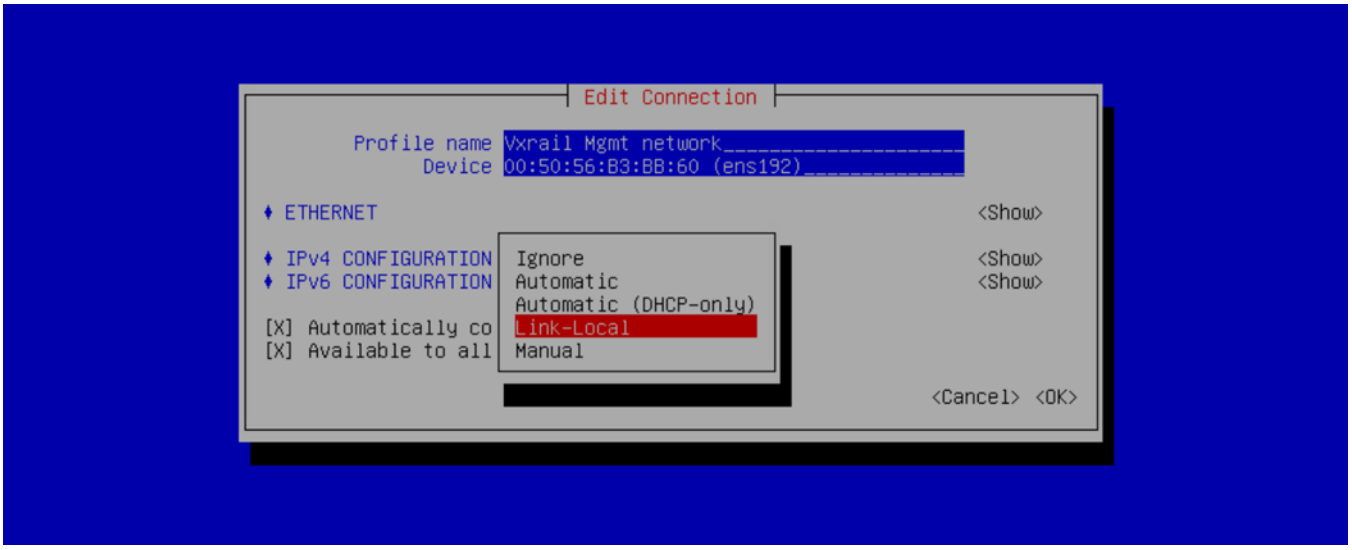

5. Click **Edit** to the right of Routing, and click **Add**.

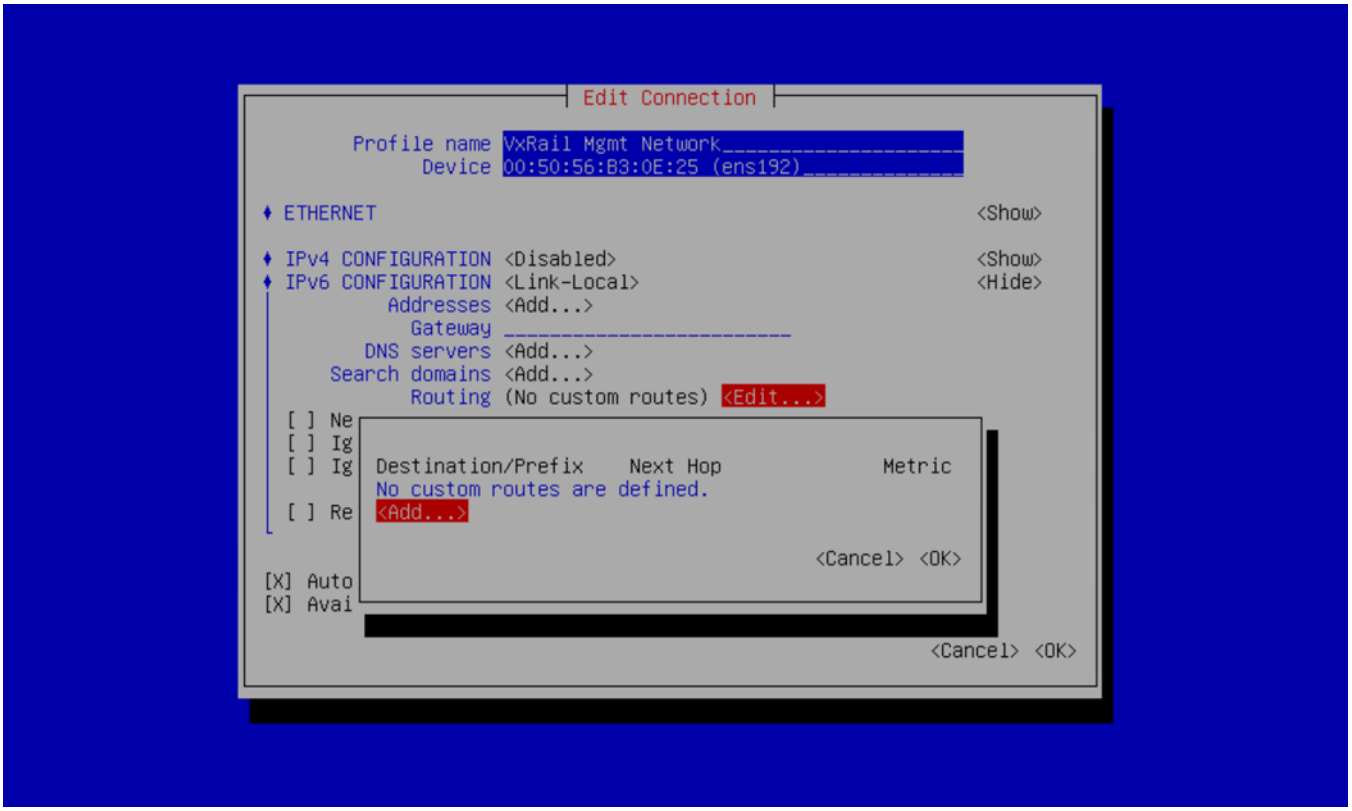

6. Enter the custom route as **fde1:53ba:e9a0:cccc::/64**, and click **OK**.

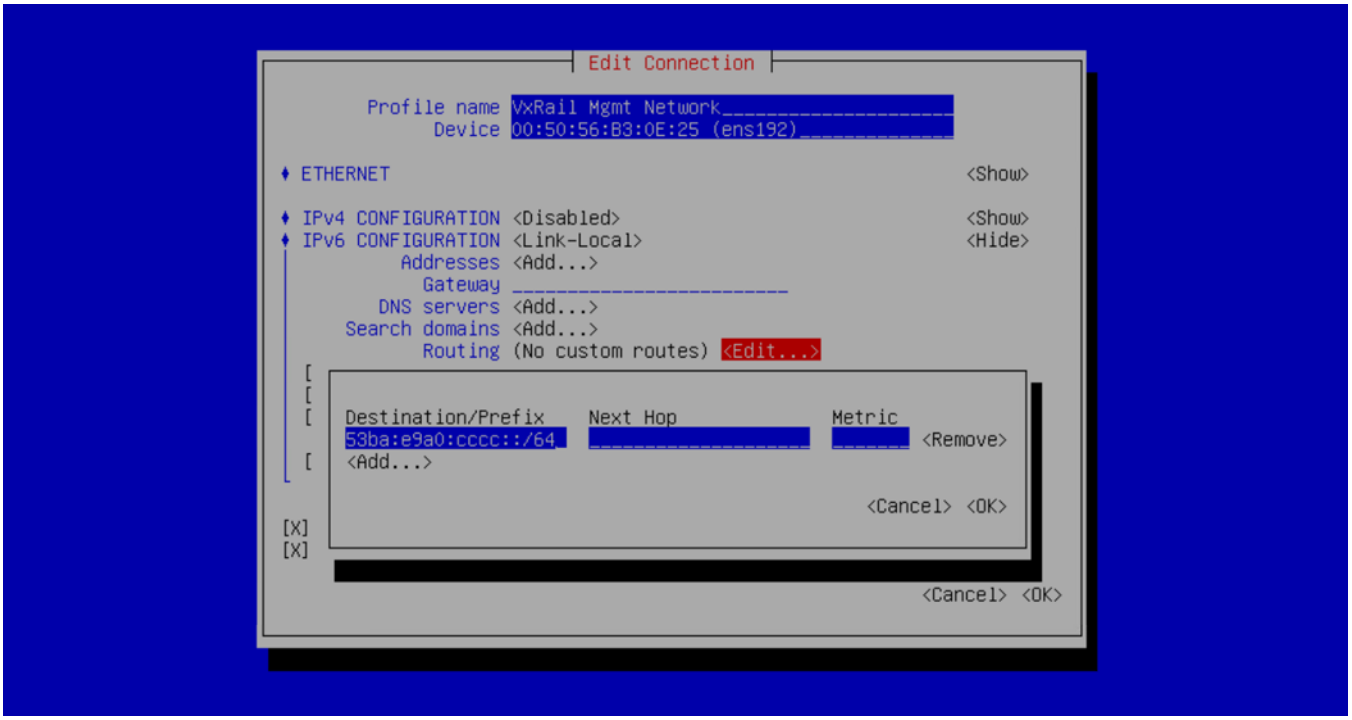

7. One custom route is now configured, click **OK**.

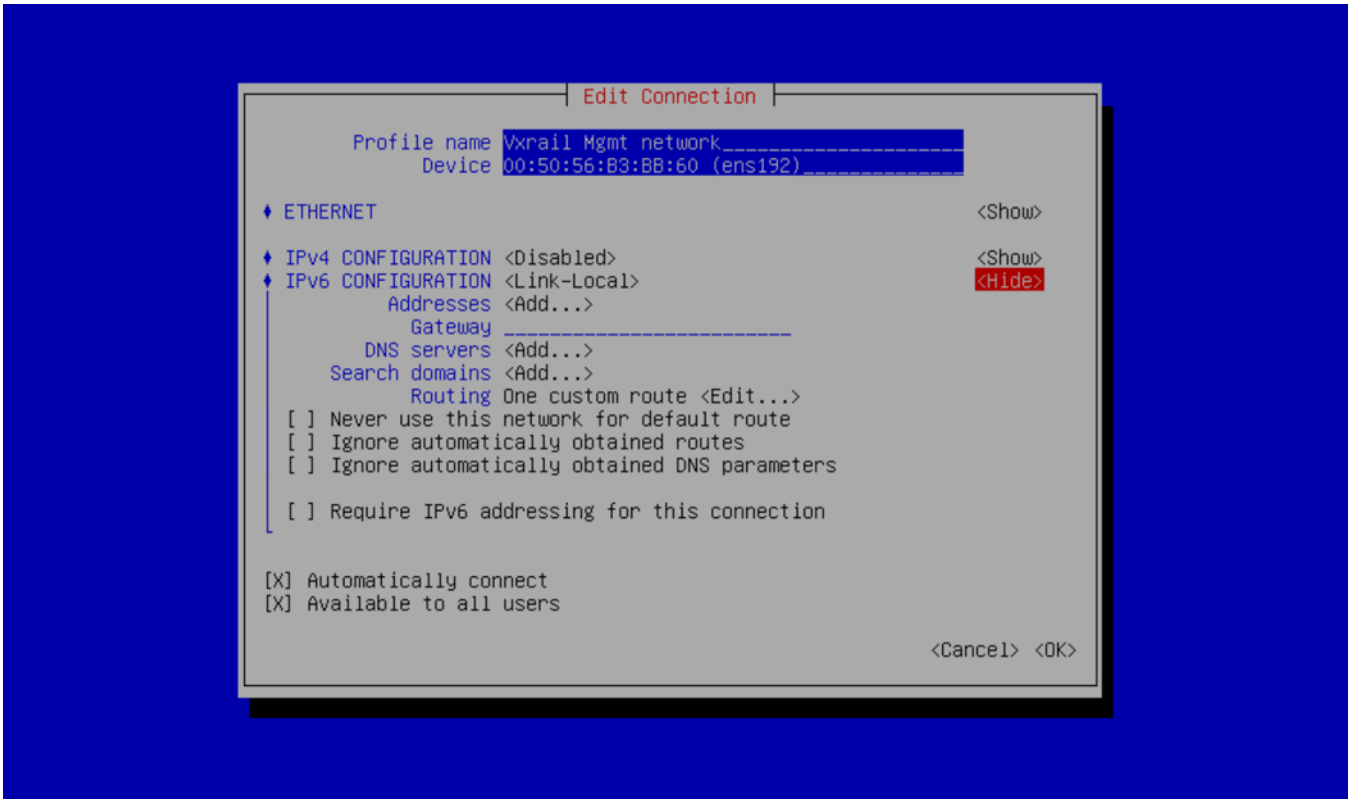

8. Click **Back** to activate the connection profiles.

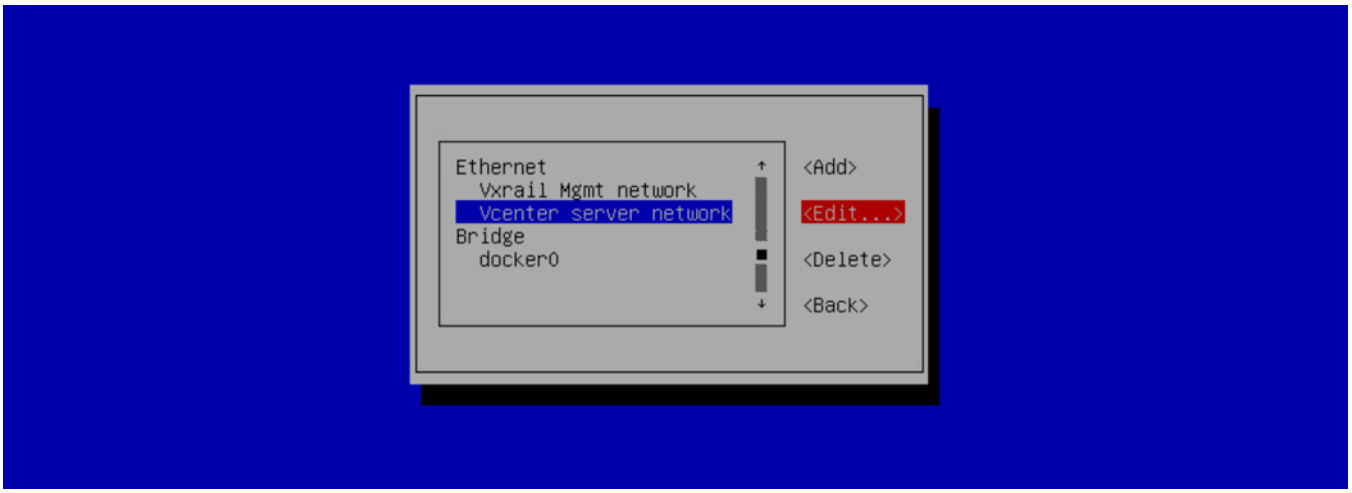

Activate connection profiles

NOTE: To populate DNS entries automatically, deactivate and active each profile.

1. Select **Activate a Connection**, and click **OK**.

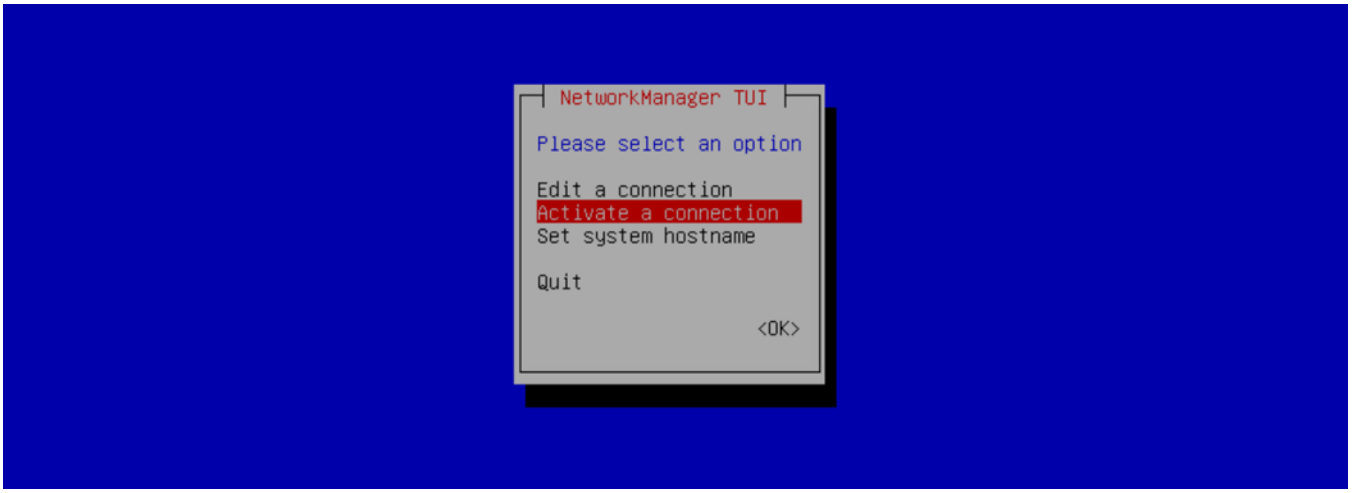

NOTE: If you change while editing a connection, you must deactivate then activate the connection for the respective interface profile.

2. Select the **vCenter Server Network** profile, and click **Deactivate**. Repeat for **VxRail Mgmt Network**.

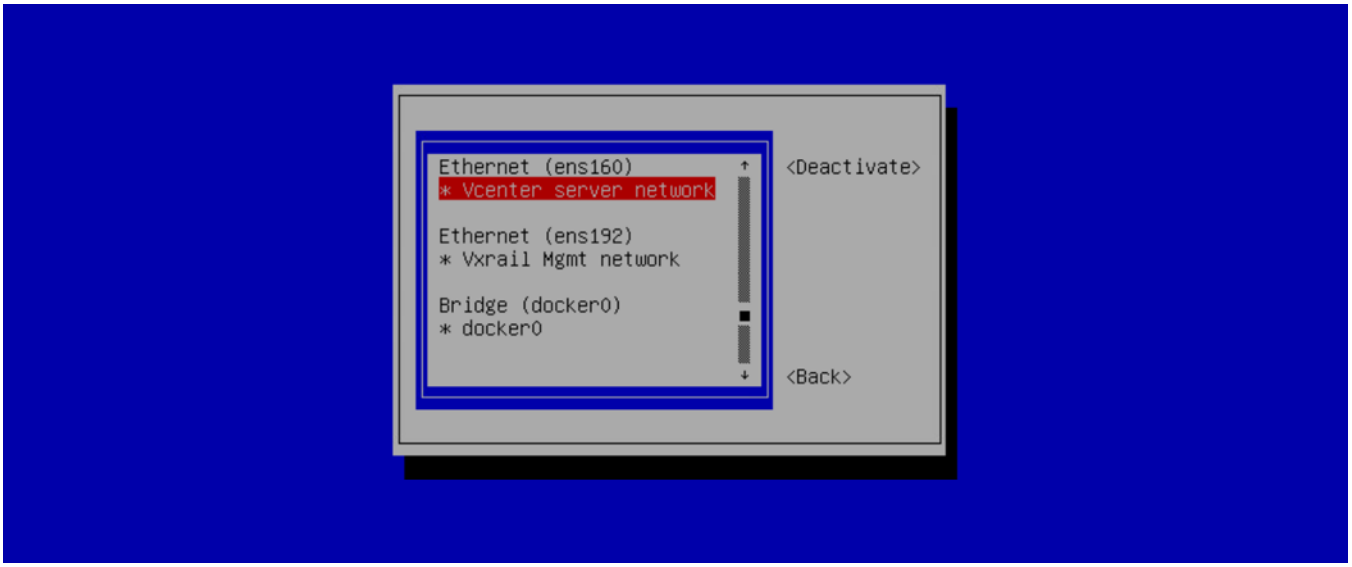

3. Select the **vCenter Server Network** profile, and click **Activate**. Repeat for **VxRail Mgmt Network**.

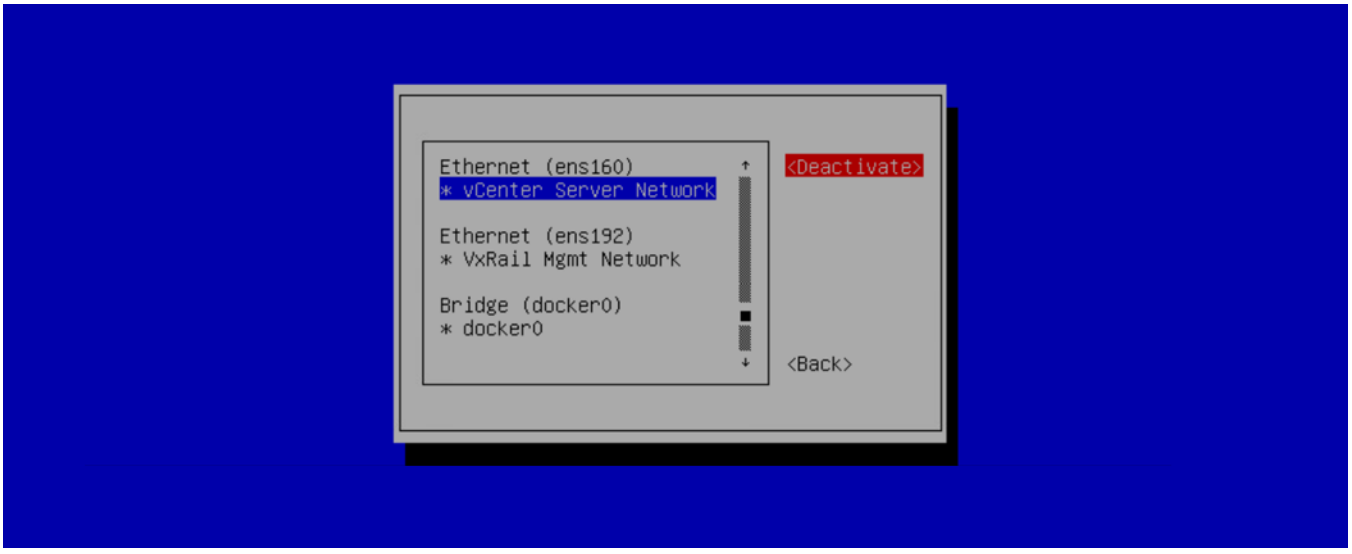

4. Click **Back**, select **Set system hostname**, and click **OK**.

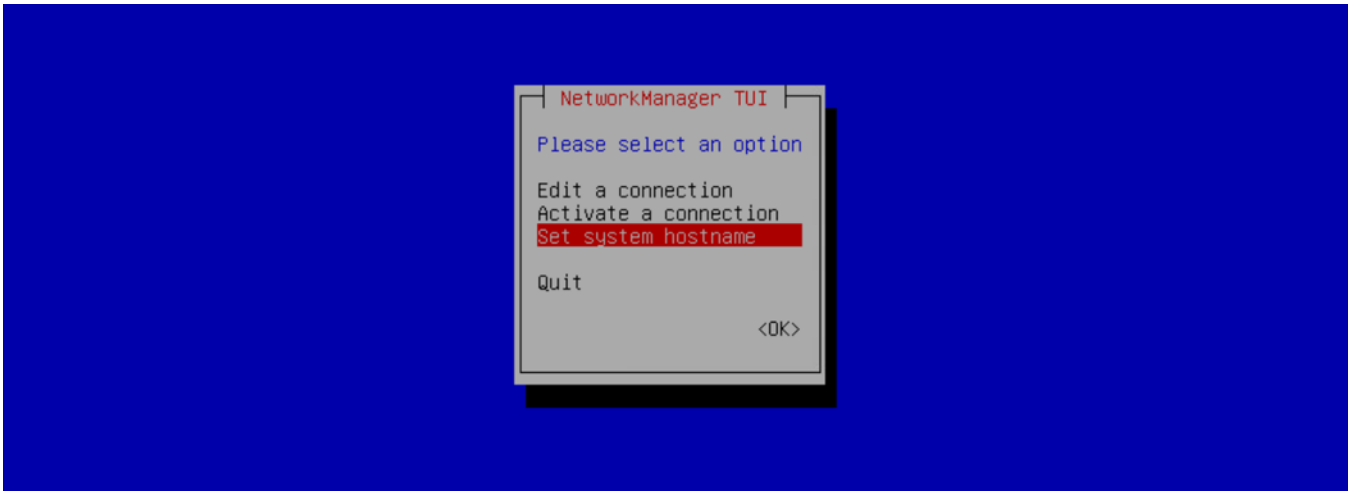

NOTE: If you are setting the hostname of OMNI, ensure you have the DNS entry of the OMNI hostname.

5. Enter **omni** for the hostname, and click **OK**.

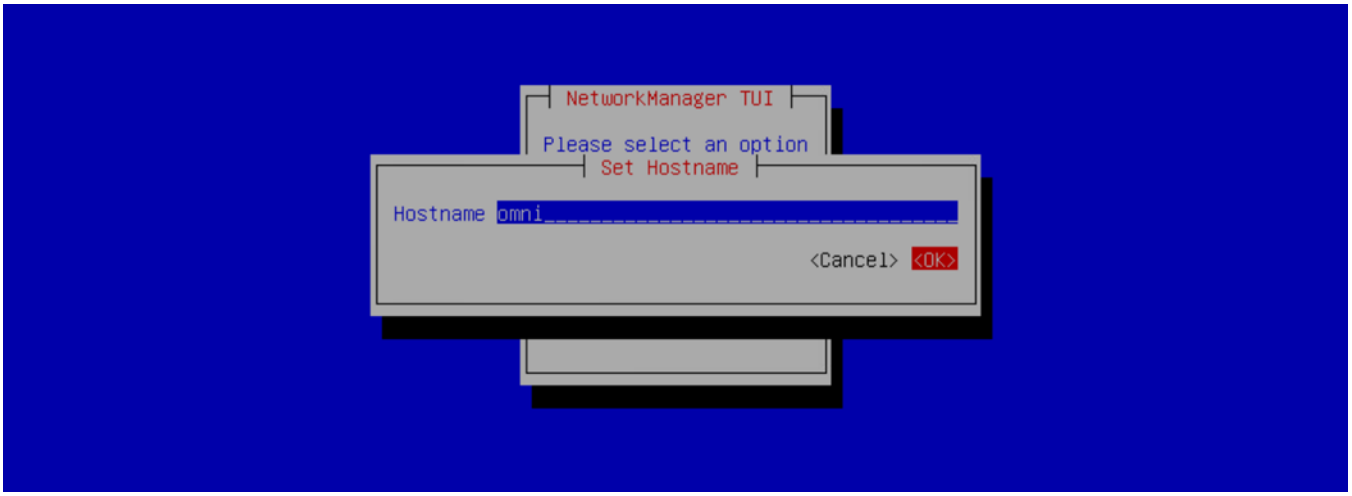

6. The hostname is now set. Click **OK**.

7. Click **Back**, and **OK** to exit the network management UI.

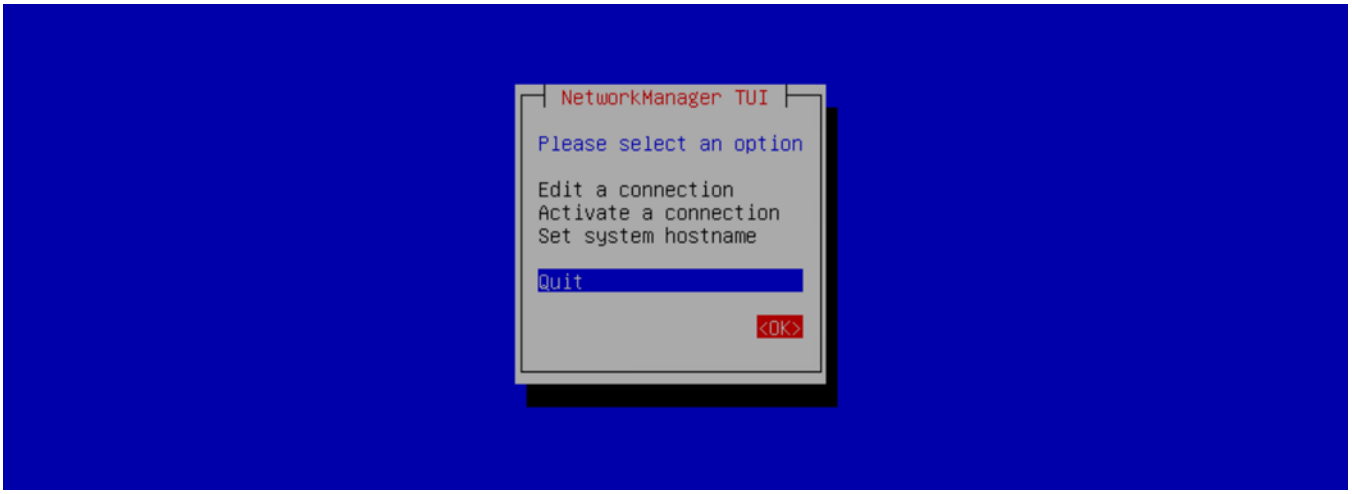

8. Enter a valid NTP Server IP address or hostname, and click **Enter**.
9. Enter **n** to not install the SSL certificate from remote server. When you enter **n**, the self-signed certificate that is created locally is installed.

   NOTE: To install a new certificate, see Generate and install SSL certificate.

NOTE: If the NTP Server is not configured, the OMNI appliance VM synchronizes with the ESXi server time zone.

Install OMNI application on ESXi server without vCenter

Starting from 2.0 release, you can install a remote OMNI instance on ESXi server without vCenter. Use this feature when you want to install OMNI independently (without vCenter), and manage SFS. For example, use this feature to install OMNI to manage SFS in Isilon deployment.

This information describes how to deploy the OMNI appliance on VMware ESXi server using the OMNI OVA file.

Prerequisite

- The supported version of the ESXi server is 6.7 or later.
- ESXi server should have the expected hardware profile to install OMNI .ova file, see OpenManage Network Integration.
- Use Chrome or Mozilla Firefox browser.

Download and Install OVA on ESXi server

1. Download the OVA from OpenManage Network Integration support, and store the OVA image locally.

2. Log in to the ESXi server.

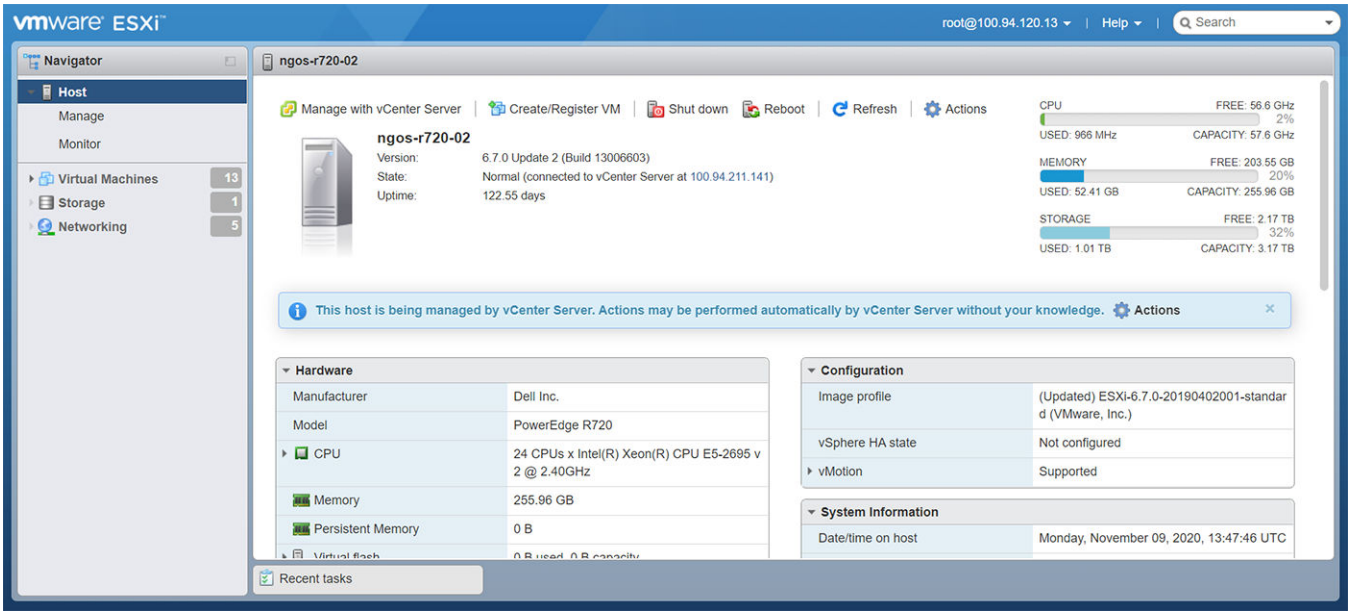

3. In the ESXi server, select **Virtual Machines**, and click **Create / Register VM**.

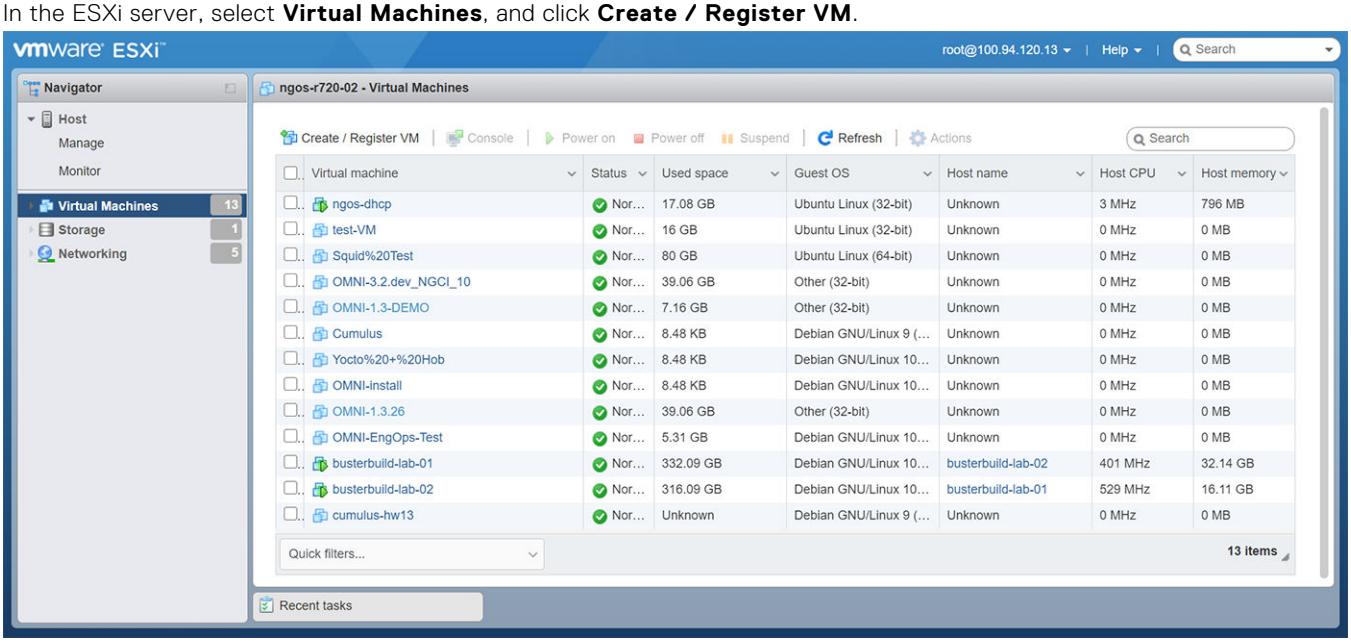

4. Select the creation type as **Deploy a VM from an OVF or OVA file** and click **Next**.

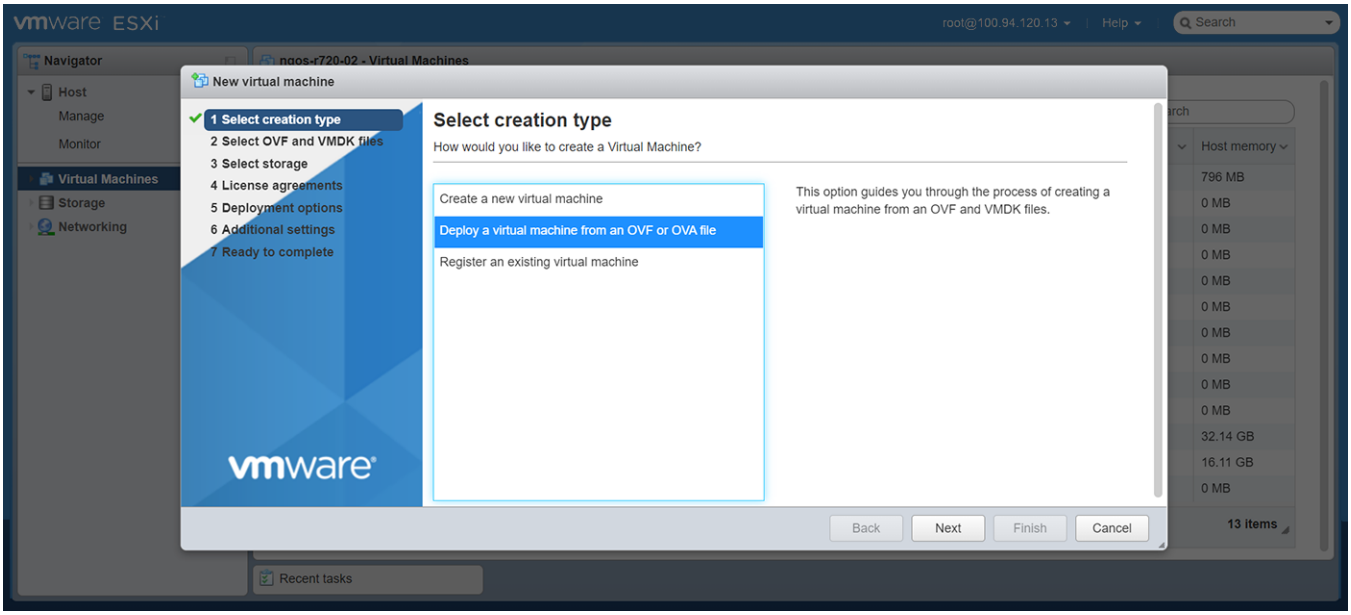

5. Enter a name for the VM and upload the OVA file from a local source, and click **Next**.

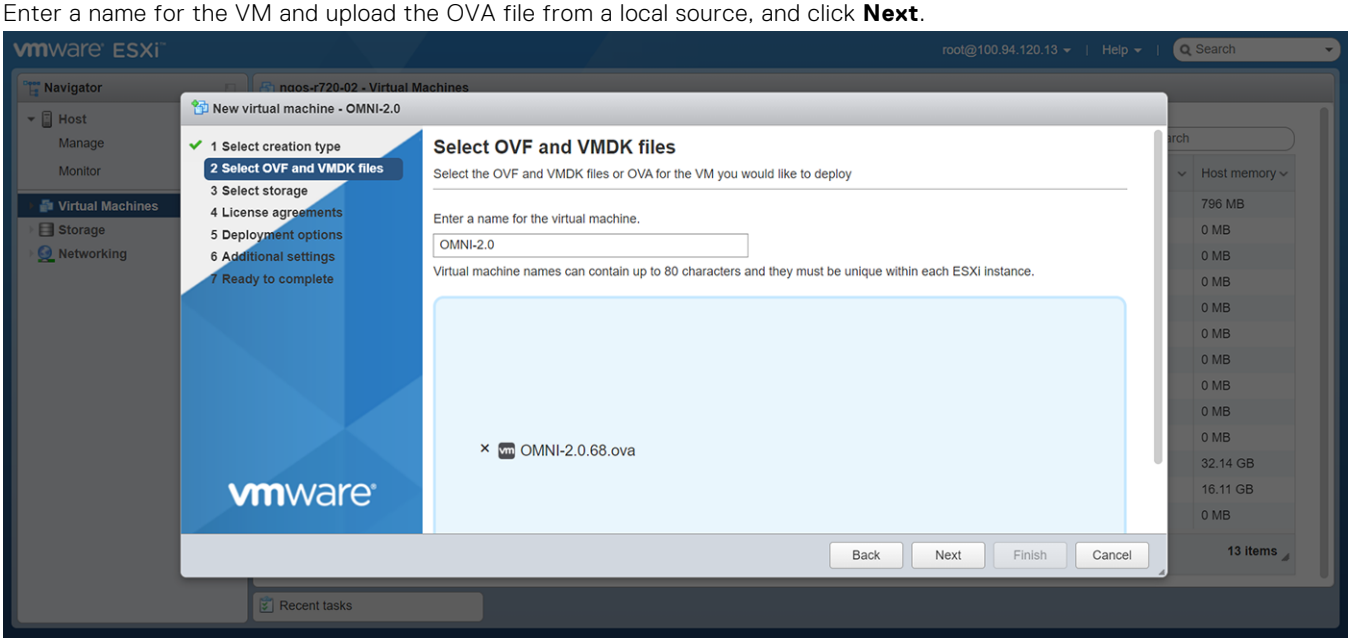

6. Select storage for VM configuration files and virtual disks and click **Next**.

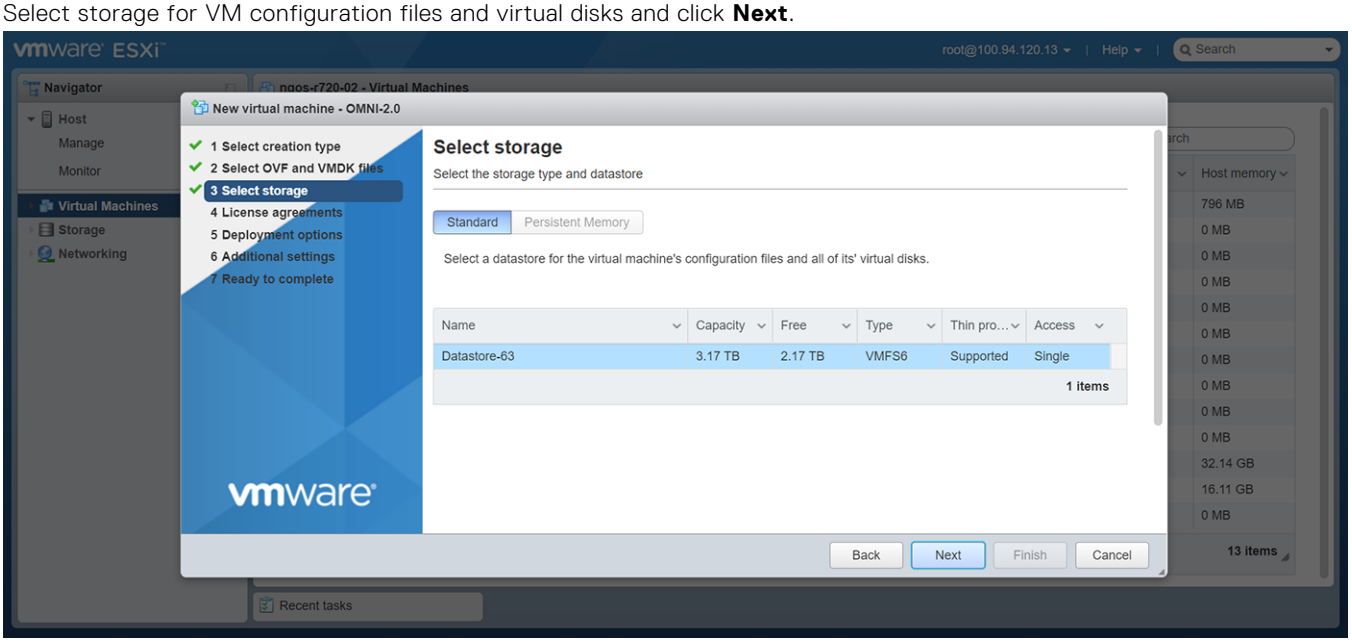

7. Accept the EULA license agreement and click **Next**.

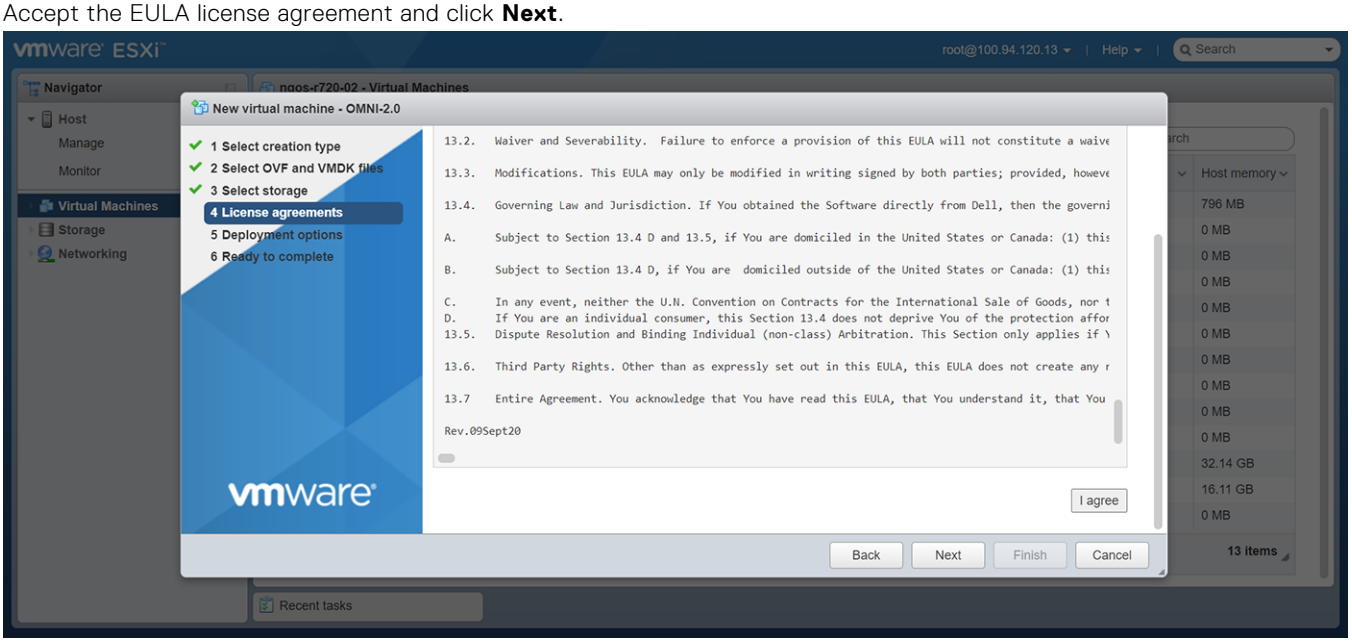

8. By default, OMNI VM OVA has dual NIC adapters. Use only one network if you deploy OMNI VM as an independent entity without vCenter. You can disable VxRail management network (ens192). Select the disk provisioning and power options, and click **Next**.

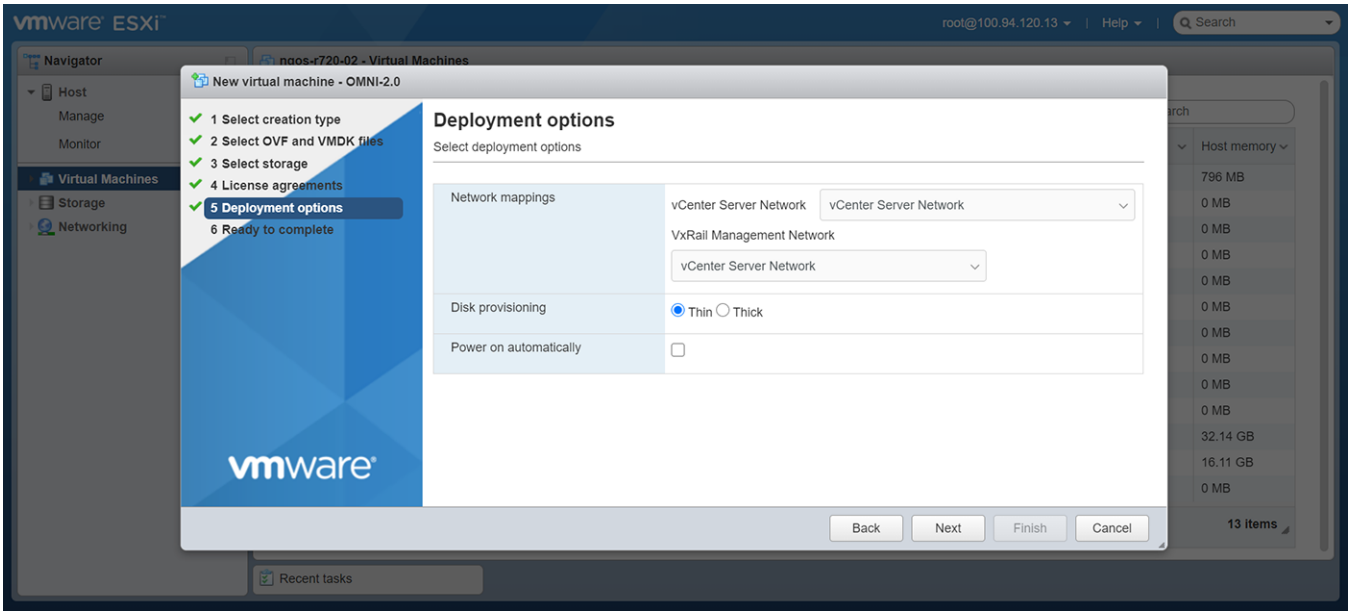

You can select **Power on automatically** checkbox to power on the VM after the installation.

9. Ready to complete page displays the summary of the settings that are configured so far. Review and verify the settings and click **Finish** to complete the installation.

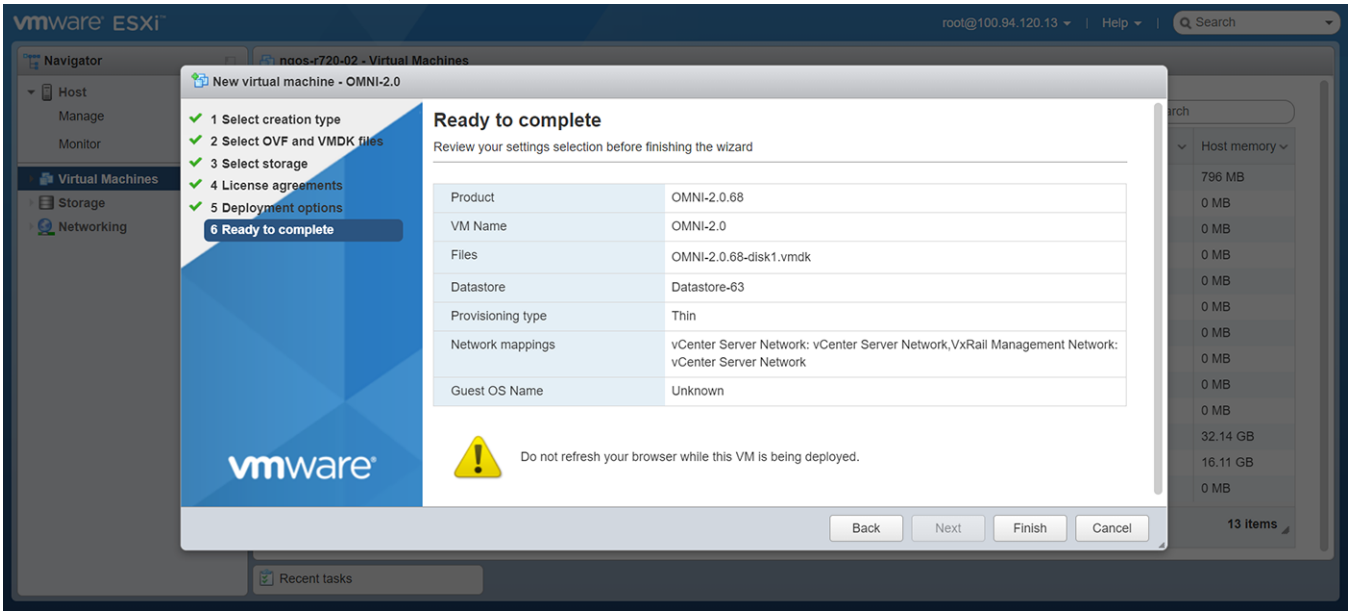

Set up OMNI

To log in to the VM console and set up OMNI configurations:

1. Follow the steps provided in the section Set up OMNI to log in to the VM.
2. Configure **Wired connection 1** (ens160) interface with the ESXi management IP address. Set the IPv4 configuration from Automatic to Manual from the drop-down and enter the required IP address details along with the subnet ask and gateway information. Set the IPv6 configuration for the interface to Ignore. See ens160 interface configuration steps from Set up OMNI.

   NOTE: Wired connection 2 (ens192) interface setup is not required for non-VxRail deployment.
3. By default, ens160 interface is activated. If you change while editing a connection, you must deactivate then activate the connection for the ens160 interface, see Set up OMNI.

OMNI appliance console CLI menu

This information describes the menus available to the admin SSH user through the console.

Table 8. OMNI appliance console CLI menu

Menu option	Submenu option	Description
1. Show version	—	Display OMNI virtual appliance and plug-in version.
2. Interface configuration menu	0. Config Docker Private network	Display default OMNI docker private network information. Also configure docker private network information. **NOTE:** OMNI default docker private subnet is 172.16.0.1/16. Dell Technologies recommends using 172.16.0.1/24 (x.x.x.x/24).
	1. Show interfaces	Display OMNI network interface configuration.
	2. Show connection status	Display OMNI network interface connection status.
	3. Configure interfaces	Configure OMNI network interfaces using Network Manager user interface (nmtui) including OMNI Management IP, gateway, DNS entries, search domains, routes, OMNI hostname, and so on.
	4. Show NTP status	Display OMNI network time protocol (NTP) server status.
	5. Configure NTP server	Configure OMNI NTP server. Enter remote NTP server IP or hostname. **It is recommended that you use the server hostname**.
	6. Unconfigure NTP server	Unconfigure OMNI NTP server.
	7. Start NTP server	Start OMNI NTP service, and enable NTP service.
	8. Stop NTP server	Stop OMNI NTP service.
	9. Exit	—
3. OMNI management service menu	1. Start OMNI management service	Start OMNI web and database essential services.
	2. View OMNI management service	Display status of OMNI essential services.
	3. Stop OMNI management service	Stop OMNI essential services.
	4. Restart OMNI management service	Restart OMNI essential services.
	5. Create support bundle	Create OMNI support bundle archive and save to download location. **NOTE:** Dell Technologies recommends using the OMNI appliance management user interface to generate and download support bundle.
	6. Change application log level	Display current log-levels, and configure DEBUG or ERROR log-levels.

Table 8. OMNI appliance console CLI menu (continued)

Menu option	Submenu option	Description
		NOTE: Dell Technologies recommends using the OMNI appliance management user interface to change log level of needed services.
	7. Exit	—
4. Password or SSL configuration	1. Change appliance password	Change appliance `admin` user password.
	2. Change root password	Assign password of application root user; root user is disabled by default, and is required to set the password first to access the root user. Root user is only accessible using the vCenter OMNI VM console. **CAUTION: Changing the system state from the Linux shell can result in undesired and unpredictable system behavior. Only use Linux shell commands to display system state and variables, or as instructed by Dell EMC Support.**
	3. Generate self-signed SSL certificates.	Replace existing OMNI appliance self-sign certificate. **NOTE:** After SSL certificate installation completes, you must re-register OMNI with the vCenter.
	4. Install SSL certificates from remote server.	Replace OMNI certificates with the certificate that is on the remote server using SCP or FTP. **NOTE:** After SSL certificate installation completes, you must re-register OMNI with the vCenter.
	5. Exit	—
5. Upgrade appliance	—	Upgrade the OMNI appliance.
6. Reboot appliance	—	Reboot the OMNI appliance.
7. Show EULA	—	Display the OMNI end user license agreement (EULA).
8. Logout	—	Log out as the `admin` user.

Generate and install SSL certificate

OMNI Management menu has options to generate self-signed SSL certificates or install SSL certificates from remote server.

Generate self-signed SSL certificate

To generate a self-signed SSL certificate:

1. From the OMNI management menu, select **4. Password/SSL configuration menu**.

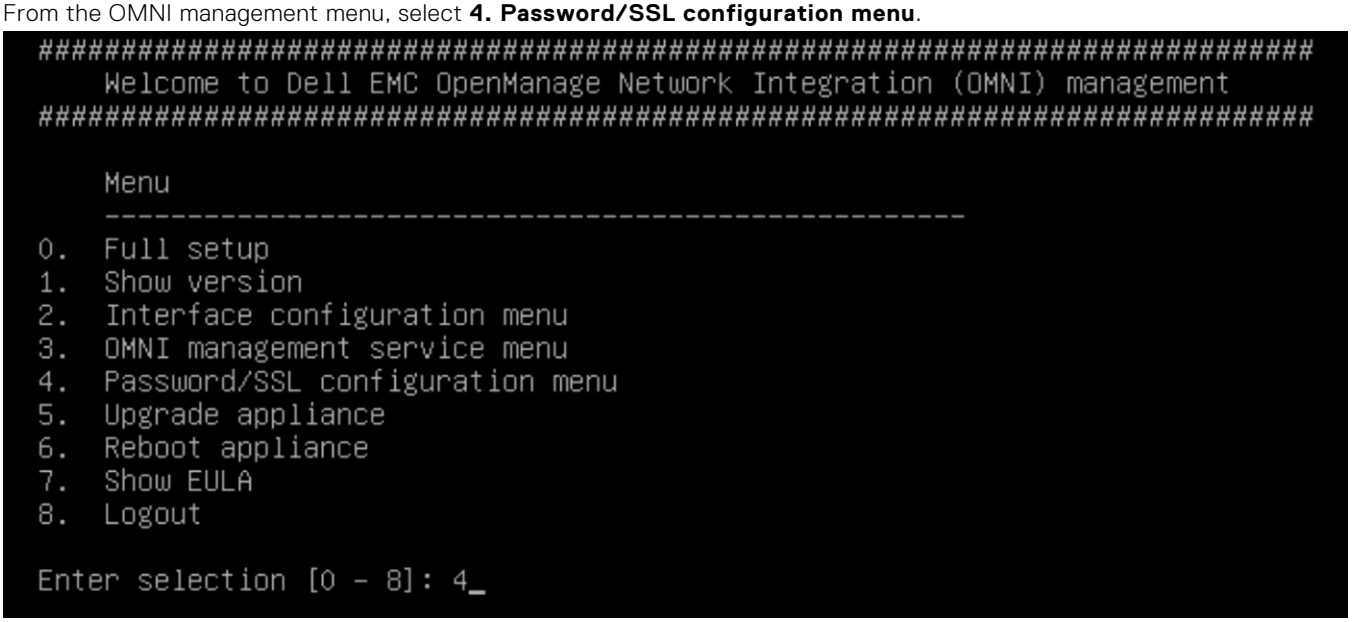

2. Enter the selection as **3. Generate self signed SSL certificates**. OMNI VM displays confirmation for replacing the existing certificate and key with the newly created certificates and keys.

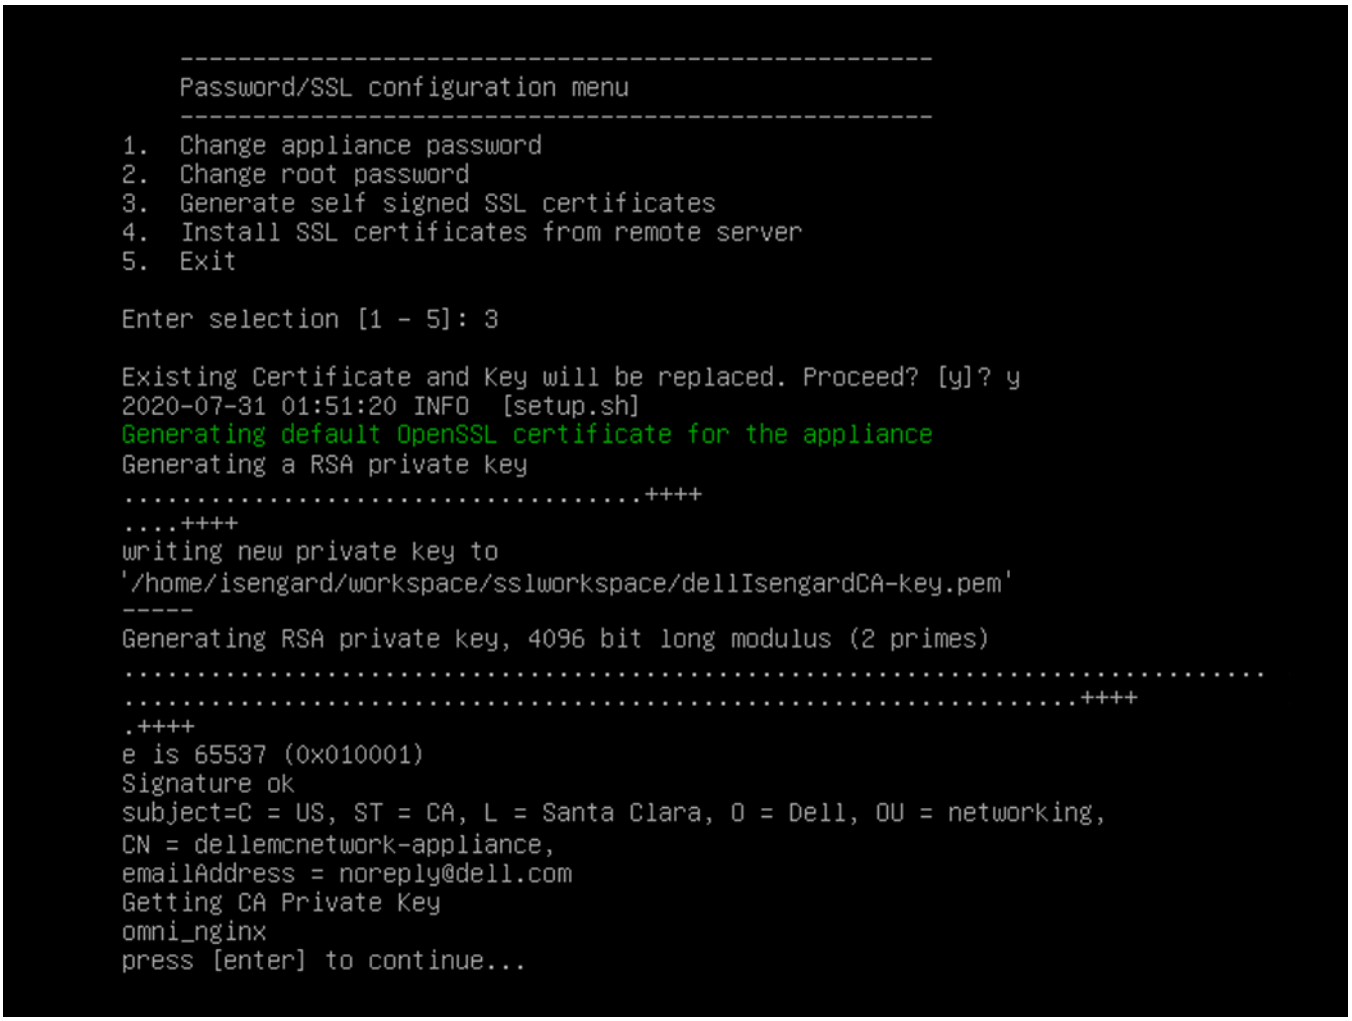

3. Unregister and register the vCenter again using OMNI UI to apply the new SSL certificate.

NOTE: Refresh the browser to view the OMNI UI plug-in from the vCenter when you register or unregister OMNI appliance with vCenter 7.0. For older versions of vCenter, log out and log in to access the plug-in from the vCenter.

Install SSL certificate from remote server

To install SSL certificate from remote server:

1. Generate SSL certificate using a standard method in .pem or .crt formats.
2. Copy the generated files to the remote SCP server.
3. From the OMNI management menu, select **4. Password/SSL configuration menu**.

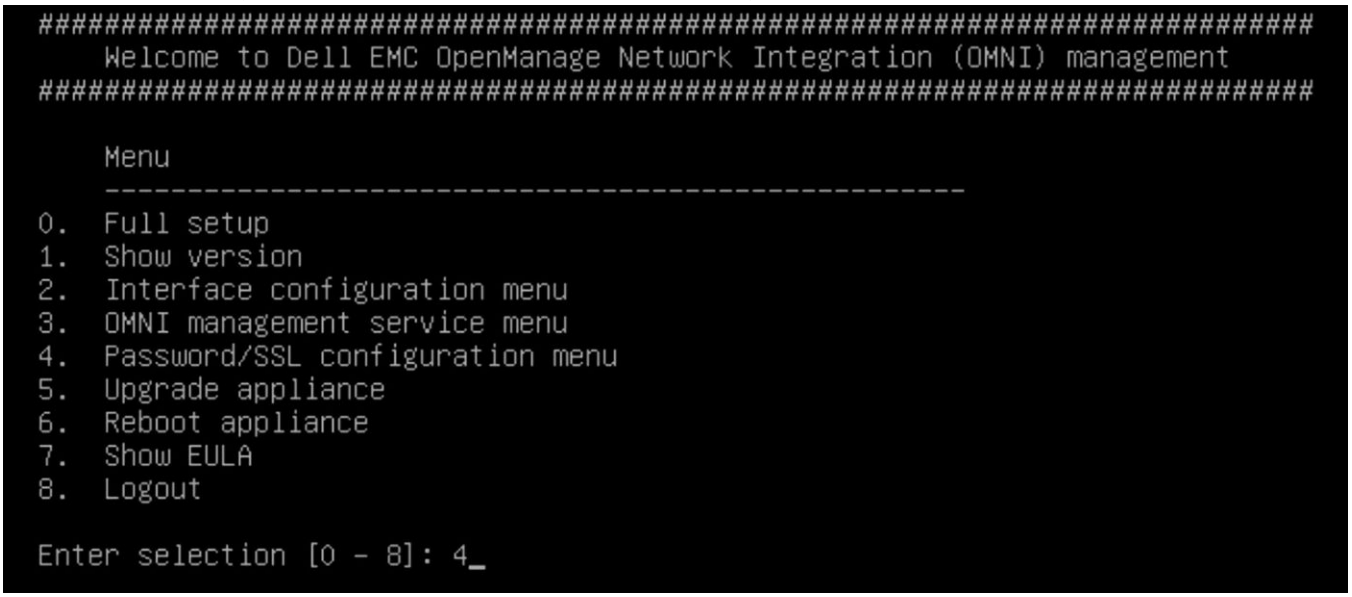

4. Enter the selection as **4. Install SSL certificate from remote server** to install the certificate.
   Enter the remote SCP server IP address or hostname and login to the SCP server. Provide the

path to the certificate and private key in the server. The files are copied to the OMNI VM.

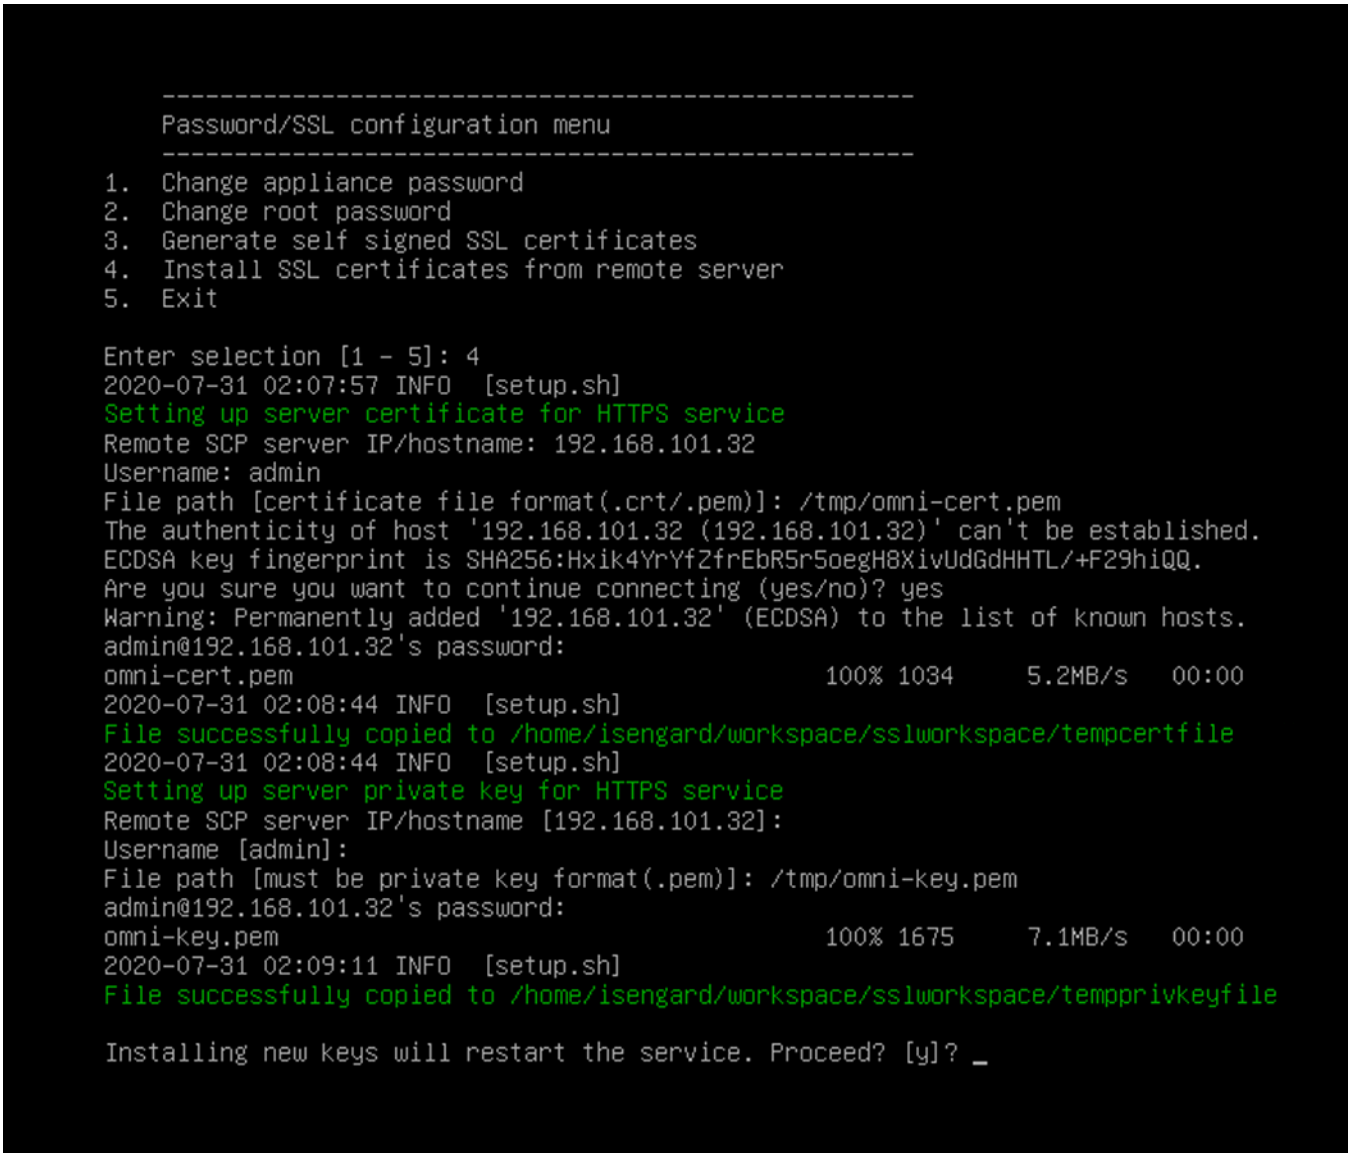

5. Enter **y** to install the SSL certificate.
6. Unregister and register the vCenter again using OMNI UI to apply the newly installed SSL certificate.

NOTE: Refresh the browser to view the OMNI UI plug-in from the vCenter when you register or unregister OMNI appliance with vCenter 7.0. For older versions of vCenter, log out and log in to access the plug-in from the vCenter.

View and configure docker private network settings

The internal docker system of the OMNI VM uses a private network to communicate with the docker components. By default, the docker bridge private network IP address is set to 172.16.0.1/16. View and change the default configuration for the docker private network using OMNI console.

View docker private network configuration

1. Log in to OMNI console.

2. From the OMNI management menu, select **2. Interface Configuration Menu**.

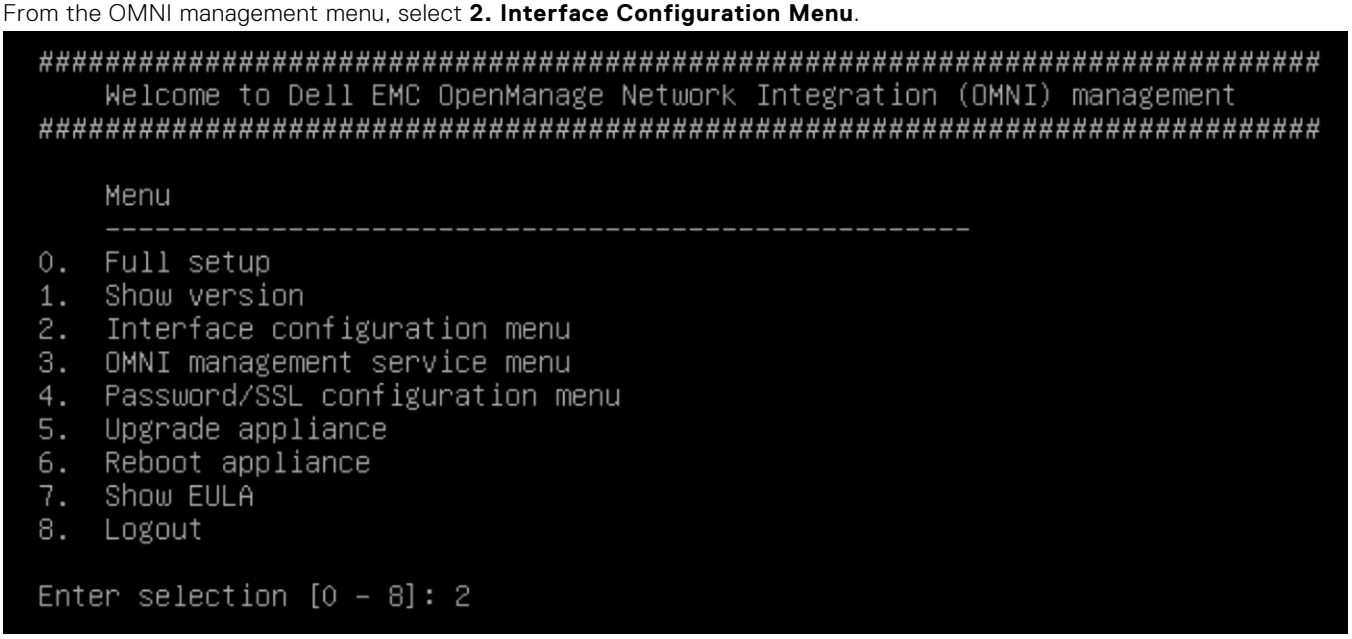

3. Select **0. Config Docker Private network**.

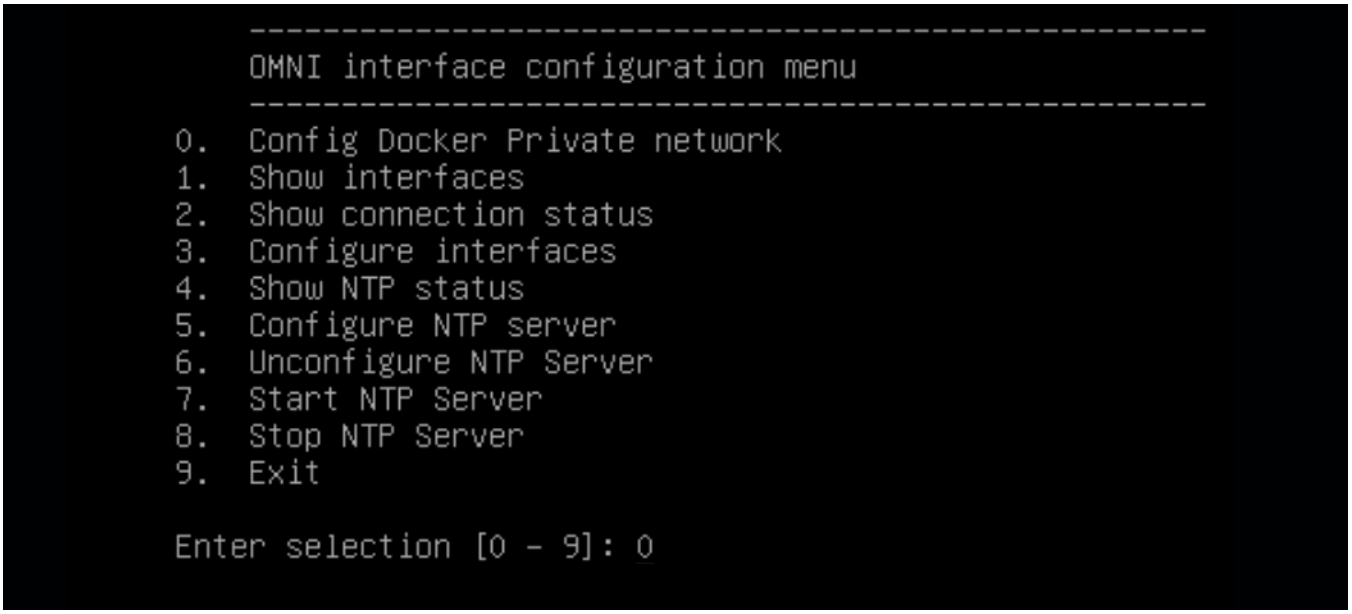

4. OMNI console displays the current docker private subnet settings with an option to change the docker private network setting. You can ignore to change the setting by entering **n**.

Change docker private network default settings

When there is a conflict between the default docker private network and any other network to which OMNI is connected, OMNI cannot communicate with the devices in that network. To avoid the conflict, you can change the docker private network default settings in OMNI.

To change the docker private network configuration:

1. From the OMNI management menu, select **2. Interface Configuration Menu**.
2. Select **0. Config Docker Private network**.
3. OMNI console displays the current docker private subnet settings. Any change to the docker private network setting results in reboot of OMNI. OMNI displays confirmation to change the docker private network. Enter **y** to proceed with the configuration change.

4. Enter the private network IPv4 network address for docker in *A.B.C.D* format with subnet mask in prefix-length */xx* format.

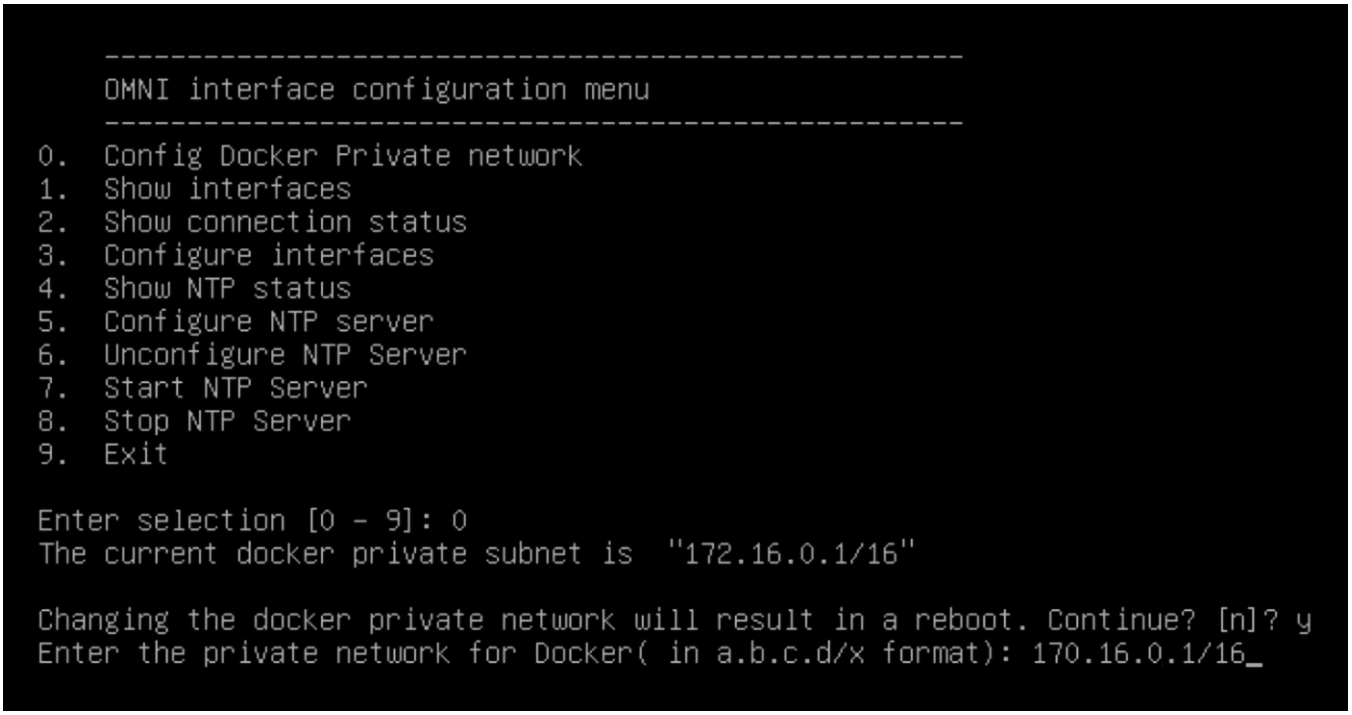

5. OMNI reboots automatically to implement the latest docker private network configuration.

5

OMNI vCenter integration

This information describes the OMNI vCenter integration to automate vCenter `PortGroup` VLANs.

vCenter VSS and DVS PortGroups

When you configure `PortGroups` of a virtual standard switch (VSS) with VLANs and distributed virtual switch (DVS) with VLANs on the OMNI registered vCenter, the respective active and standby physical adapter interfaces are automatically configured by OMNI on the SmartFabric `ServerInterfaces`. This is shown as tasks on the registered vCenter tasks pane.

CAUTION: If you remove all port groups from the host, OMNI discovers that no port group is assigned to the host and deletes all the networks.

DVS provides an option to change the VLAN of uplink `PortGroups`. OMNI ignores `PortGroup` configuration if the VLAN type `PortGroup` is set to VLAN trunking or private VLAN.

Dell Technologies recommends keeping the DVS uplink in Trunking mode and configures the virtual `PortGroups` with VLANs for each network. OMNI configures the respective VLANs on the ToRs and SmartFabric uplinks.

OMNI automates the vCenter `PortGroup` VLAN and manages the registered vCenter by identifying the relation between the SmartFabric `ServerInterface` and the ESXi host PNIC MAC.

Identification of vCenter ESXi Host by OMNI

OMNI collects the PNIC MACs of all ESXi hosts in registered vCenters. If OMNI identifies the `ServerInterface` ID as a collected PNIC MAC (Id=MAC without ':') of the host, OMNI identifies that host to belong to an OMNI registered SmartFabric instance.

Table 9. vCenter `PortGroup` VLAN automation of identified ESXi host

vCenter action	SmartFabric action by OMNI
Add or update `PortGroup`: VLAN of VSS or DVS.	• Create network of `PortGroup` VLAN and set `Network Originator` to `Auto`. • Add network to SmartFabric `ServerInterface`.
Remove `PortGroup` from VSS or DVS.	Remove unused networks from SmartFabric `ServerInterface` .

NOTE: OMNI automation is not designed to delete unused `ServerInterfaces` of SmartFabric. If you remove ESXi host from vCenter, OMNI does not remove the networks that are associated with the server interfaces. Dell Technologies recommends removing the network that is associated with the server interface profile manually.

SmartFabric networks consolidation by OMNI

As part of automation, OMNI:

1. Collects all networks of registered SmartFabric.
2. Collects networks of `ServerInterface` of registered SmartFabric.
3. Identifies SmartFabric networks that are created by the OMNI UI, and SmartFabric networks that are not created by the OMNI UI.

   OMNI distinguishes the origin of the network that is configured through vCenter or OMNI user by setting the `Network Originator` parameter.

a. OMNI sets `Network Originator` to `Manual` when you create a network using OMNI UI.
b. OMNI sets `Network Originator` to `Auto` when OMNI vCenter PortGroup automation creates a network.

4. Appends networks that are not created by the OMNI UI (all networks except the network that has `Network Originator` set to `Manual`) to SmartFabric uplink of the type `Default` or `CreateOnly`.

   OMNI automates network addition on one of the fabric uplinks. If you edit the network on the uplink using the OMNI UI and add or edit Network of Originator type `Auto`, the automation process may remove that network.

5. Finds unused networks; SmartFabric networks that are not created by the OMNI UI, and not used by the SmartFabric `ServerInterface` and SmartFabric uplinks.
6. Deletes unused networks from the SmartFabric.

NOTE: The `Default` or `CreateOnly` uplink can be configured on the SmartFabric through the OMNI. For more information, see Configure and manage uplinks.

NOTE: If there is any change to IP address or hostname of a VxRail node in the cluster, OMNI takes about 20 to 30 minutes to reflect any new network configuration changes for the changed VxRail host.

6

Access the OMNI stand-alone portal

You can access OMNI as a stand-alone portal using the OMNI IP address. OMNI appliance page displays links to launch the OMNI Appliance Management UI, OMNI Fabric Management Portal, and OMNI Documentation. You can access the OMNI UI using the latest version of the browsers, such as:

- Google Chrome
- Mozilla Firefox

Starting from release 2.0, OMNI provides more secure and scalable secure sign-on feature, when launching OMNI as a stand-alone user interface:

- **Logout**—Manually terminate the login session using the **Log out** button in the upper right of the UI.
- **Login session timeout**—OMNI terminates an inactive login session after 15 minutes to prevent unauthorized access.

NOTE: This feature is not applicable if OMNI is launched from vCenter plug-in.

To access the OMNI UI as a stand-alone application:

Open a browser session, go to **https://*OMNI_IP*** with the configured IP address or FQDN.

Access the OMNI Fabric Management Portal

1. From the OMNI stand-alone page, click **Launch OMNI Appliance Management**.
2. Click **Launch OMNI Fabric Management Portal** link to go to the login page.

- Enter the **username** and **password** for the OMNI VM, then click **Sign In**.

NOTE: Alternatively, you can also log in to **Fabric Management portal** directly using **https://*OMNI_IP*/delawareos10** with the configured IP address or FQDN.

After successful authentication, **OMNI Home** page is displayed.

Once you log in to the OMNI Fabric Management Portal with the username and password, **OMNI Home** page is displayed.

As an initial step, you must configure OMNI IP address or FQDN first to proceed with any other configuration in OMNI, see Configure OMNI. After configuring OMNI IP address or FQDN, from **OMNI Home**, you can:

- Add, edit, or delete SmartFabric instance, see here.
- Add, edit, or delete NSX-T instance, see here.
- Add, edit, or delete OME-Modular, see here.
- Add, edit, or delete vCenter instance, see here.
- View relationship details.
- View OMNI information.

Configure OMNI

If you are logging in to OMNI for the first time, after reaching **OMNI Home** page, configure OMNI first. You cannot use other features of OMNI until you set the OMNI IP address or FQDN.

1. From **OMNI Home** > **OMNI Configuration**, click **Set OMNI Address**.

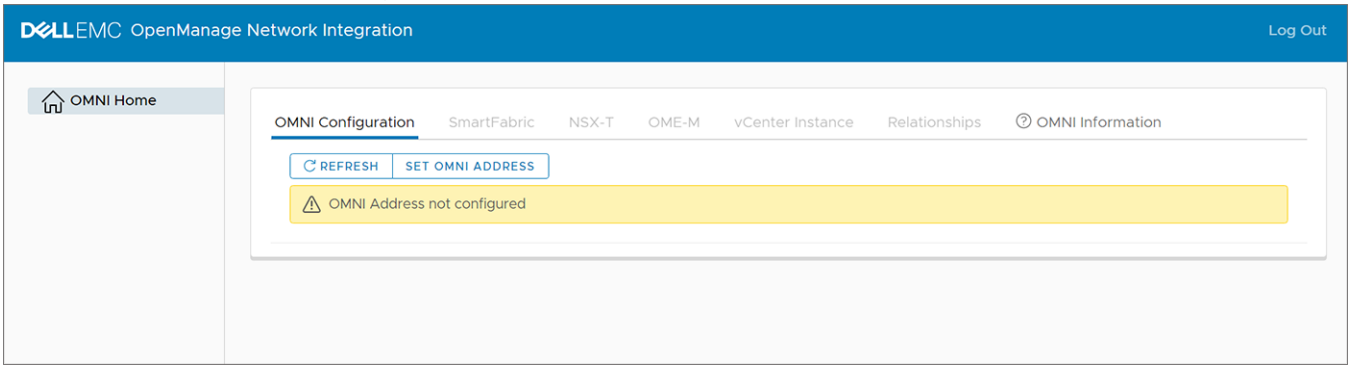

2. Enter the IPv4 address or FQDN of the OMNI and click **Submit**.

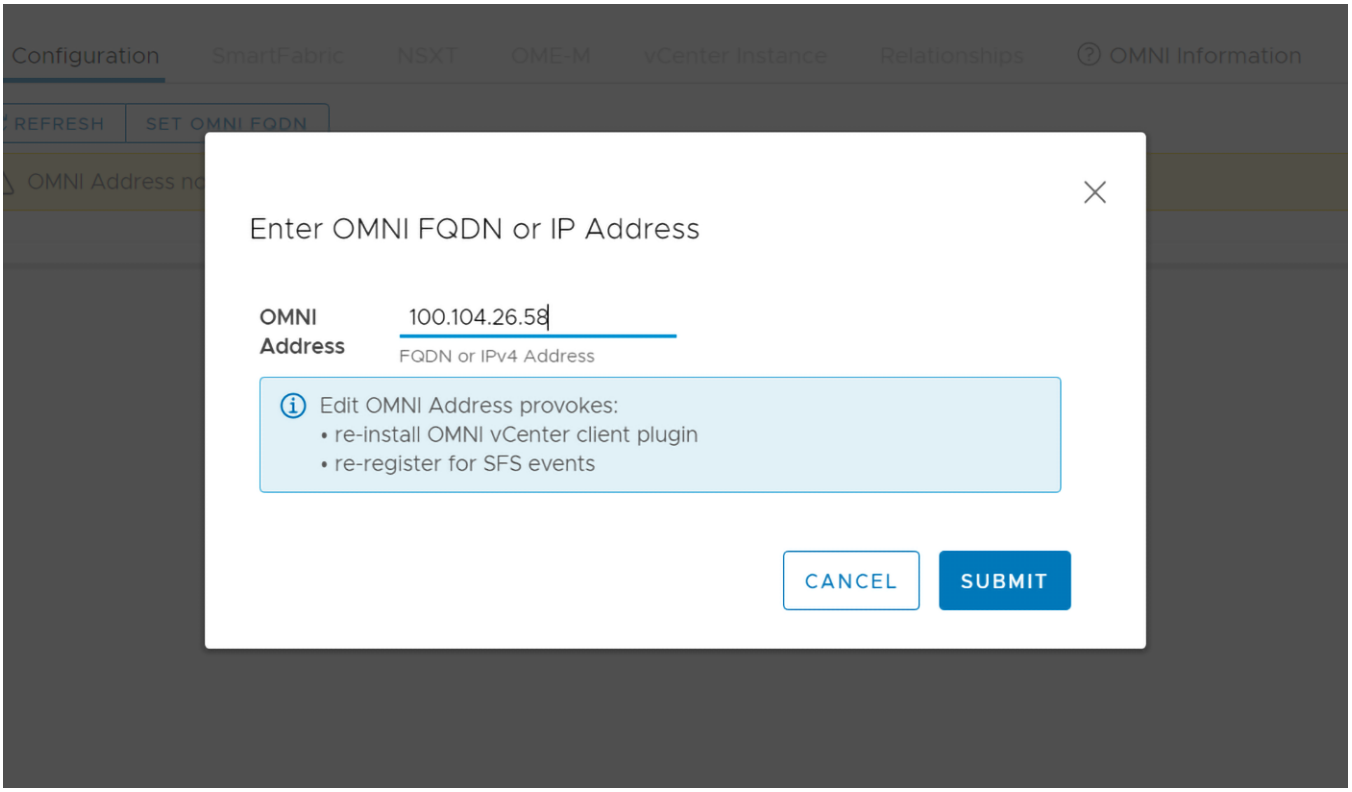

3. The system displays the OMNI address details and also enables all other options in UI.

Register vCenter with OMNI

Starting from 2.0 release, you have to register the vCenter instance with OMNI using the UI. You can register up to 8 vCenters in a single OMNI VM.

To register the vCenter with OMNI:

1. Go to **OMNI Home** > **vCenter Instance**.

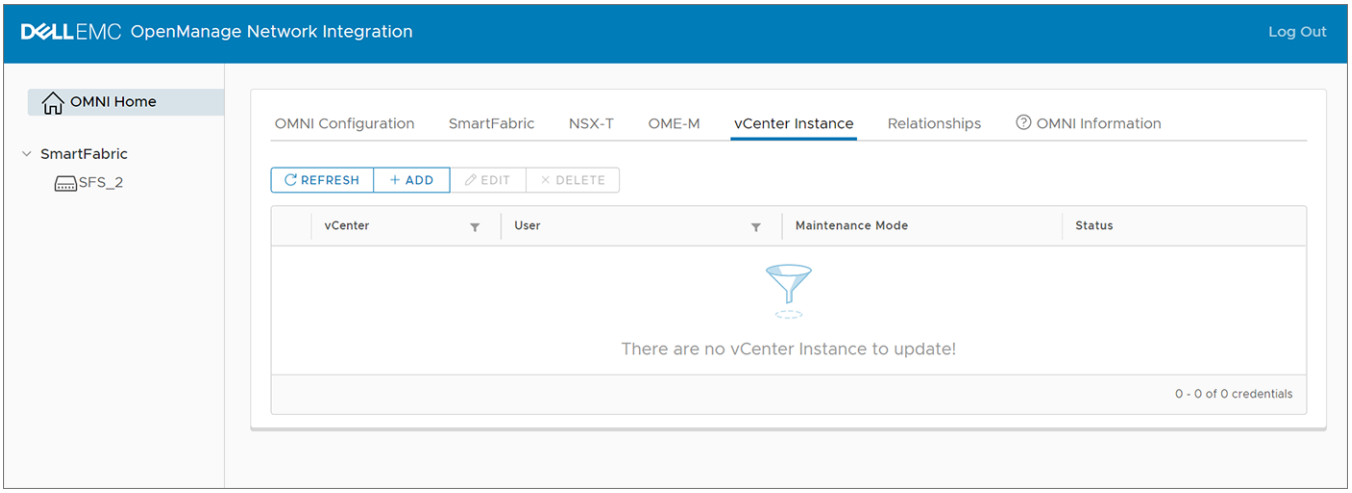

2. Click **Add** to register the vCenter.
3. Enter the IP address or FQDN of the vCenter, username, and password.
4. Select the appropriate options for **Registered** and **Automation** options. By default, **True** is selected for Registered and Automation.
   - Registered—Selecting **True** registers the vCenter with OMNI.
   - Automation—Selecting **True** creates and starts the automation service for that vCenter.

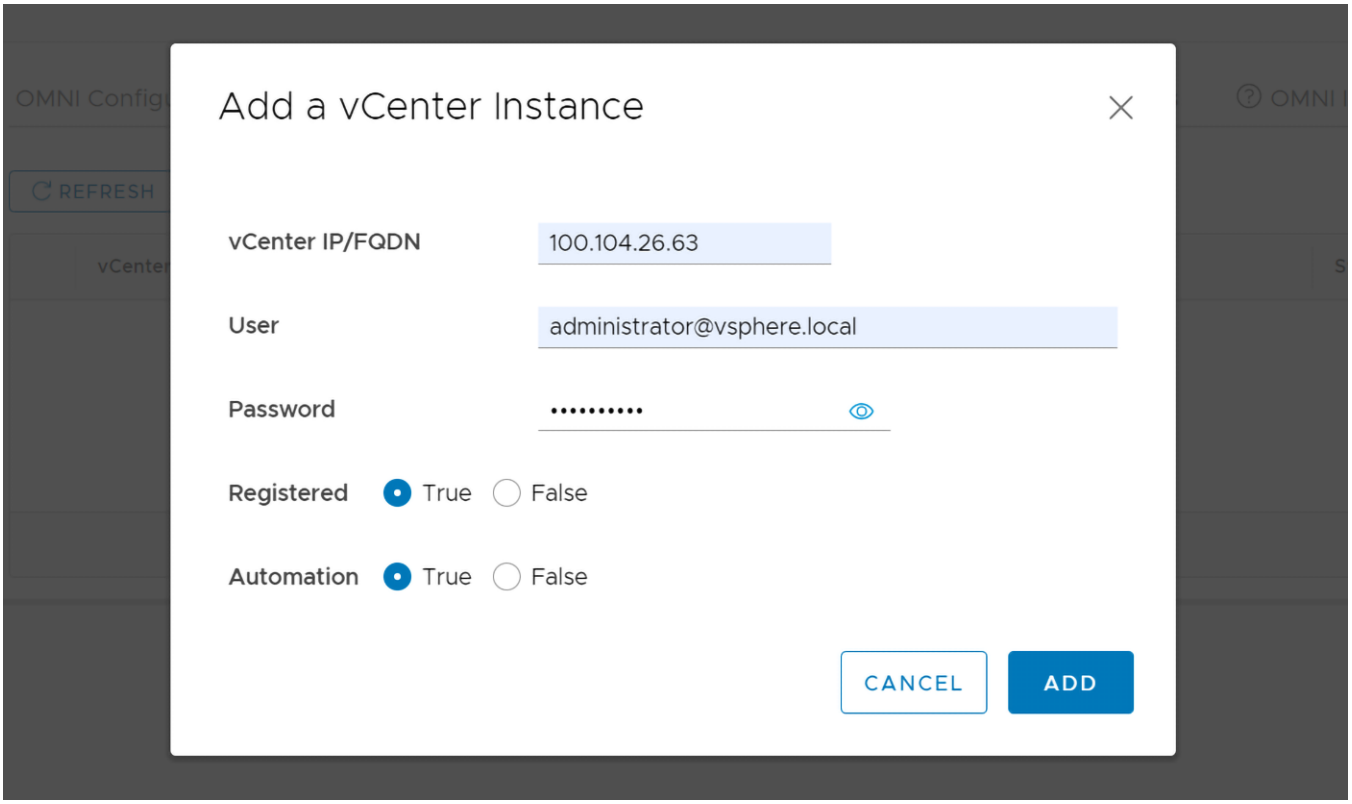

5. The system displays a vCenter registration successful message.

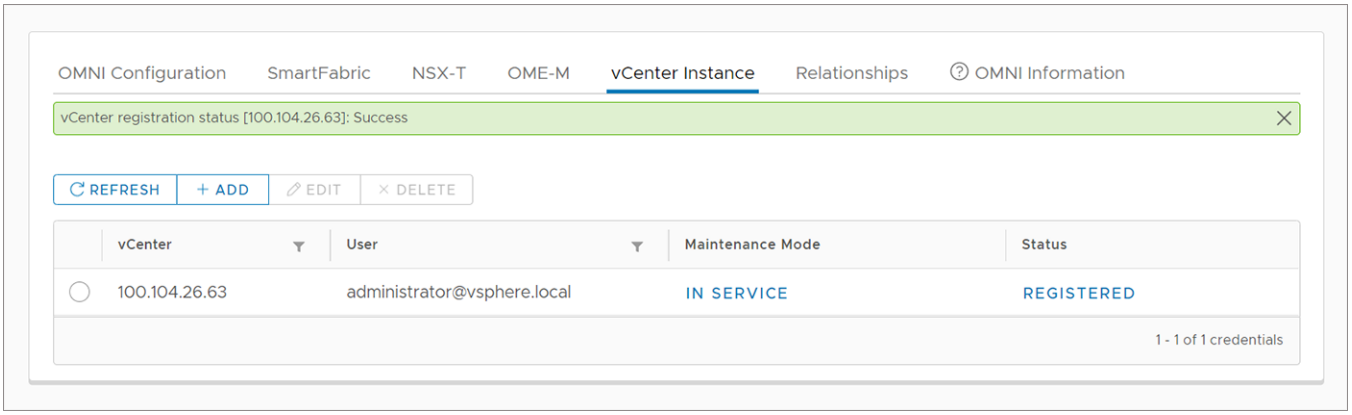

When adding the vCenter instance, you can choose only to add the instance and not register. To do so, select **False** for **Registered** option. Selecting **False** adds the vCenter and no register the vCenter with OMNI. You can register later without entering the credentials again by changing the status.

You can choose not to enable the automation services for the vCenter by selecting **False** for **Automation** option. Selecting **False** creates the automation service for the specific vCenter. You can start the automation service for the vCenter whenever required, see vCenter Maintenance mode.

Edit a vCenter instance

To edit the vCenter configuration:

1. Select the vCenter instance that you want to edit and click **Edit**.

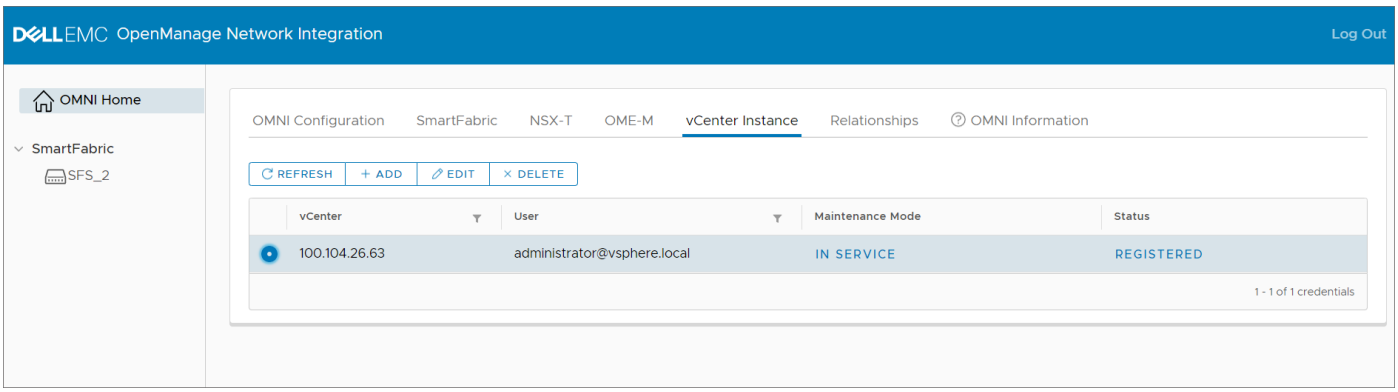

2. Update the password and click **Edit**.

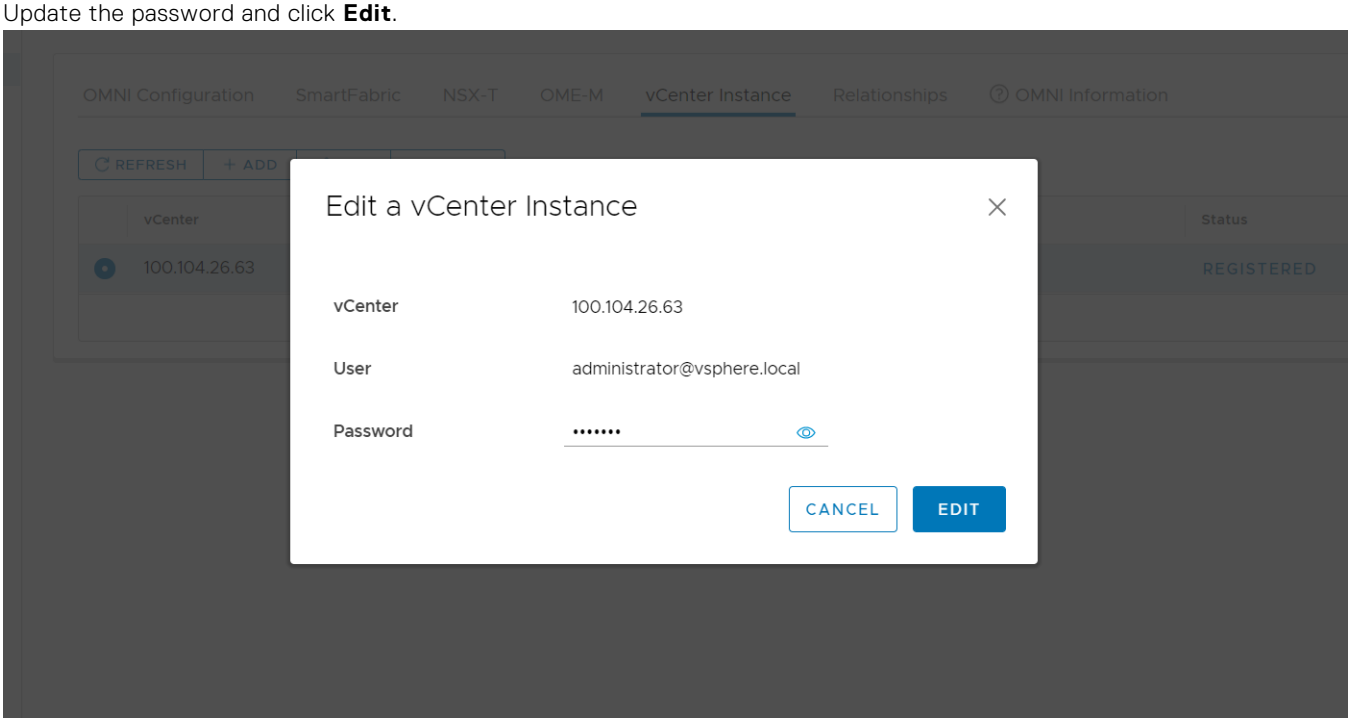

Delete a vCenter instance

To delete the vCenter configuration:

1. Select the vCenter instance that you want to delete and click **Delete**.

2. Click **Delete** to confirm the deletion.

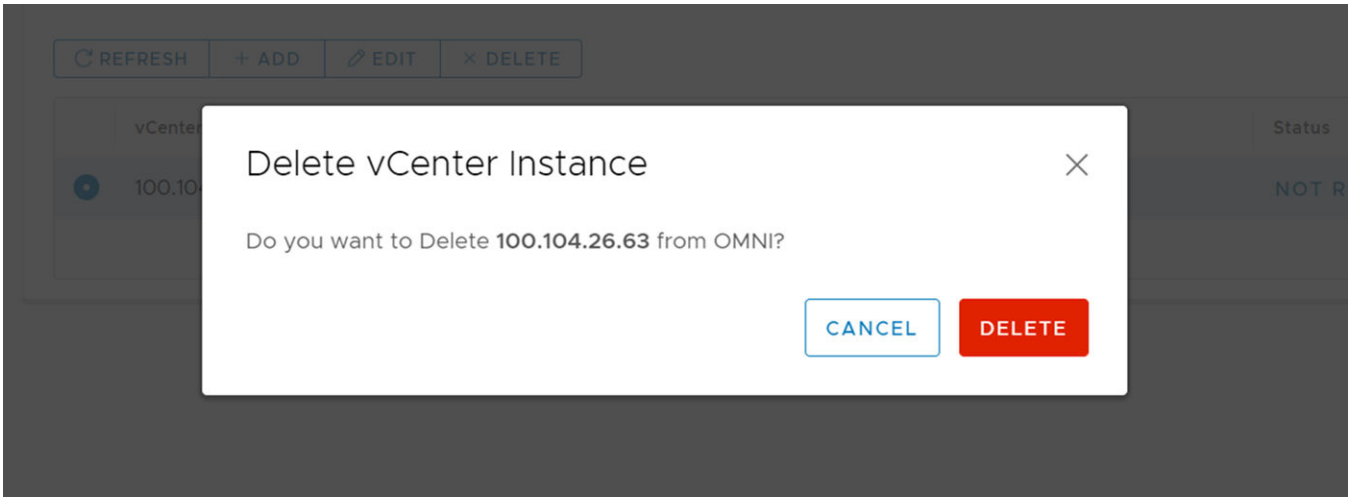

3. The system displays deletion success message.

Unregister vCenter

To unregister the OMNI from vCenter:

1. Go to **OMNI Home** > **vCenter instance**, select the status information **Registered** of the vCenter you wanted to unregister.

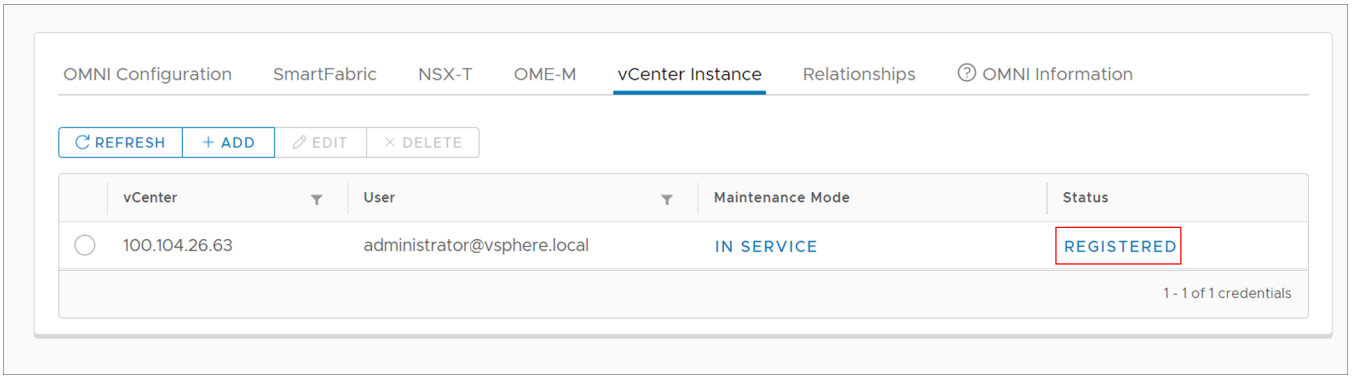

2. Click **Ok** to unregister the OMNI from the vCenter.

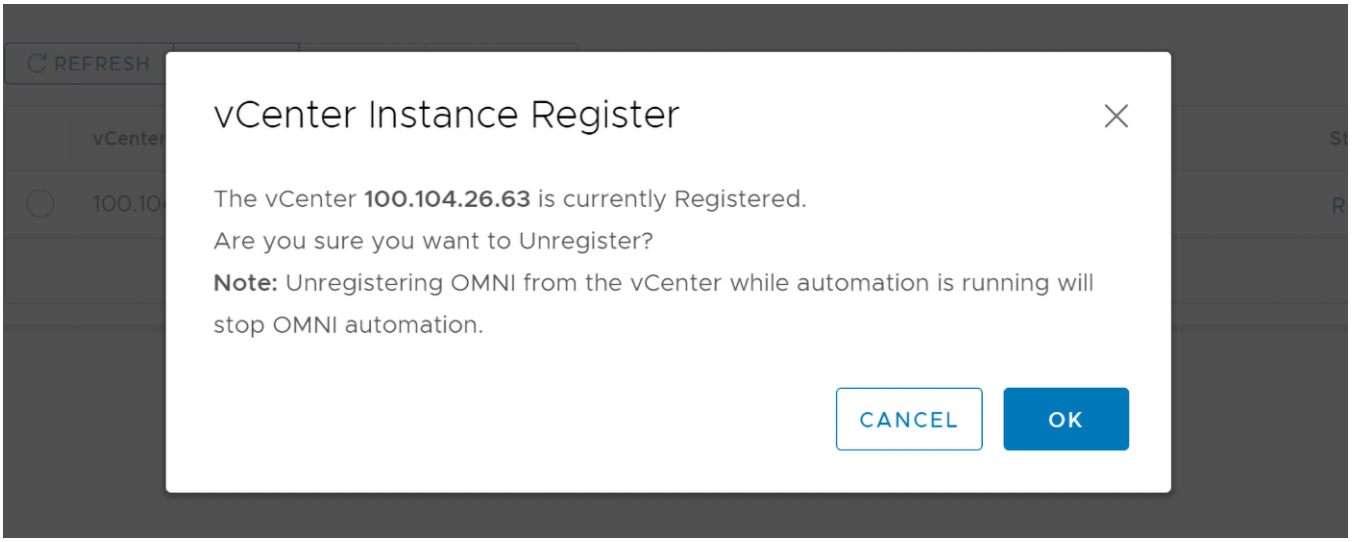

3. The system displays status change success message.

The status and the Maintenance mode for the vCenter changes to **Not Registered** and **Under Maintenance** respectively. When you unregister the vCenter, OMNI stops the automation services for that vCenter.

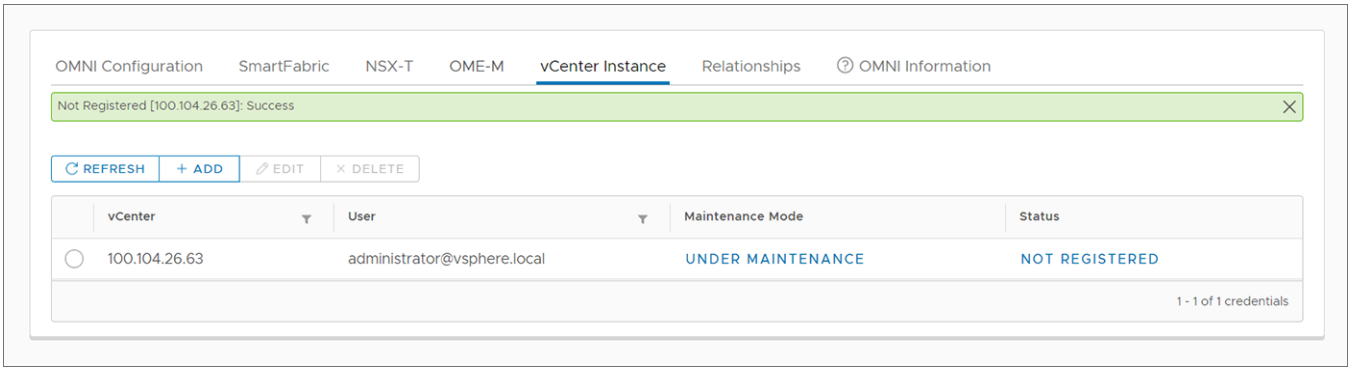

NOTE: When you unregister the vCenter, the OMNI plug-in is undeployed from vCenter. With vCenter 7.0, refresh the browser to see the change. For older versions of vCenter, log out and log in to see the changes.

vCenter Maintenance mode

Enabling Maintenance mode for vCenter instance disables automation for all SmartFabric instances that are registered with that vCenter.

Enable Maintenance mode

Changing the Maintenance mode from **In Service** to **Under Maintenance** stops the automation services that is running for that vCenter. Enable Maintenance mode for vCenter instance:

1. Go to **OMNI Home** > **vCenter Instance** and click **In Service** mode for a specific vCenter instance.

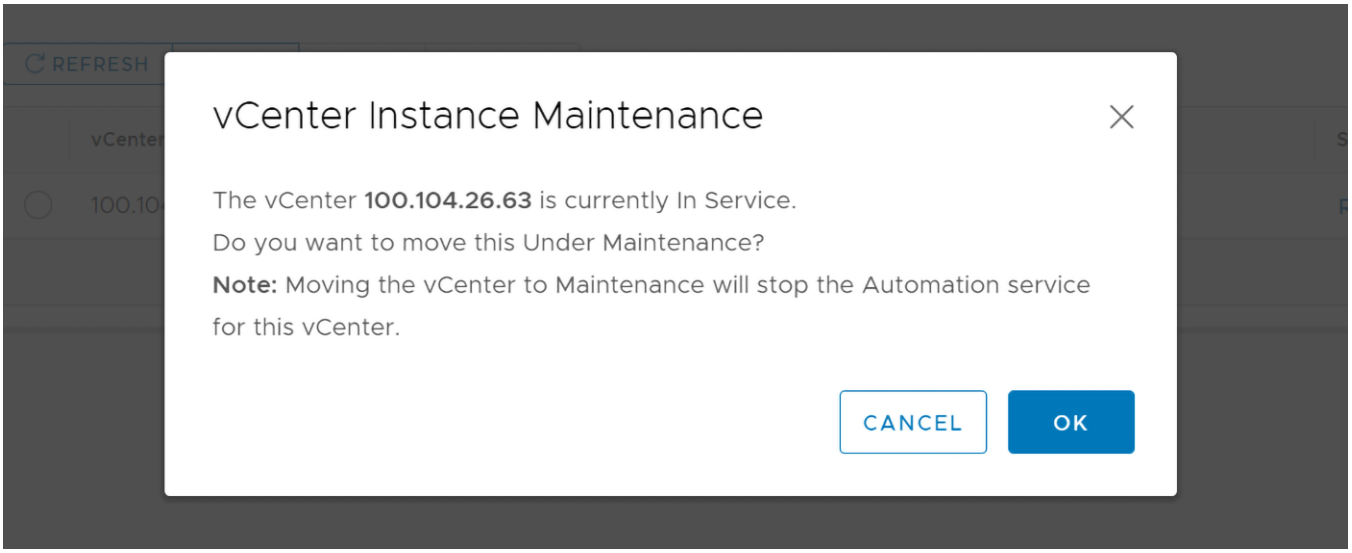

2. Click **Ok** to confirm. This action changes the mode from **In Service** to **Under Maintenance** and stops OMNI from configuring networks on SmartFabric when there are changes in the vCenter port groups through automation.
3. The system displays Maintenance mode change success message.

Disable Maintenance mode

Changing the Maintenance mode from **Under Maintenance** to **In Service** enables the vCenter to be active. To disable Maintenance Mode for a vCenter instance:

1. Go to **Home** > **vCenter Instance** and click **Under Maintenance** for a specific vCenter instance.

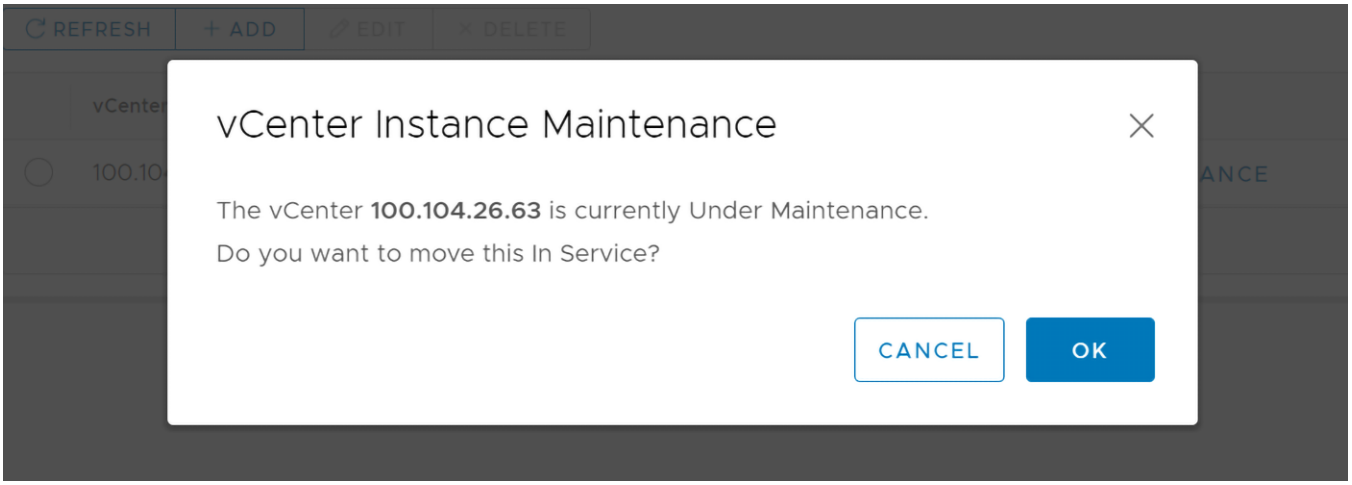

2. Click **Ok** to confirm. The vCenter status changes to **In Service** and OMNI starts the automation service for the vCenter.
3. The system displays Maintenance mode change success message.

Access OMNI plug-in from the vCenter

This information describes how to access OMNI plug-in from the vCenter. After you register vCenter with OMNI, a shortcut is available from the vSphere Client left-pane within the menu drop-down and shortcuts view.

Before you use the plug-in, you must set up an OMNI appliance in vSphere. Once you register OMNI with vCenter, the OMNI plug-in is available in the vCenter. For more information about how to register vCenter with OMNI, see here.

NOTE: vCenter 7.0 supports plug-in autodiscovery feature. So, when you register or unregister OMNI appliance with vCenter 7.0, refresh the browser to view the OMNI UI plug-in from the vCenter. When using older versions of vCenter, log out and log in to access the plug-in from the vCenter.

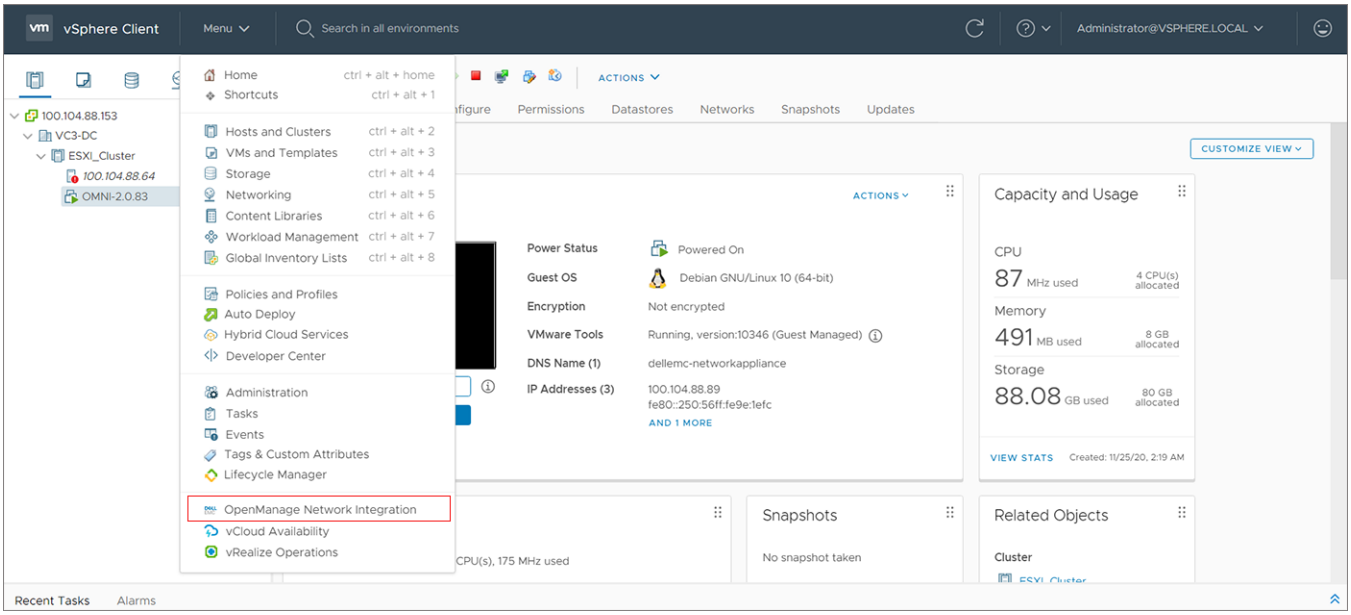

When you select **OpenManage Network Integration**, the **OMNI Home** page is launched. You can add SmartFabric, NSX-T, and OME-M instances and manage the service instances.

Edit OMNI autodiscovered SmartFabric instance

This information describes how to configure OMNI autodiscovered SmartFabric instances. If the OMNI virtual appliance is connected to a link-local network on SmartFabric (such as VxRail Management Network-VLAN 3939), the SmartFabric IPv6 VIP is autodiscovered by OMNI. For complete information about discovery, see *mDNS service* in Fabric creation.

NOTE: This configuration is applicable only for VxRail deployment and not for PowerEdge MX environment.

When you launch the OMNI plug-in from vCenter for the first time after registering, the autodiscovered SmartFabric instance is disabled. You must edit the instance and change the REST_USER password to proceed.

Edit the autodiscovered SmartFabric instance for the REST_USER password to complete the configuration.

1. Go to the OMNI portal.
2. Select the autodiscovered SmartFabric instance from the list, and click **Edit**.

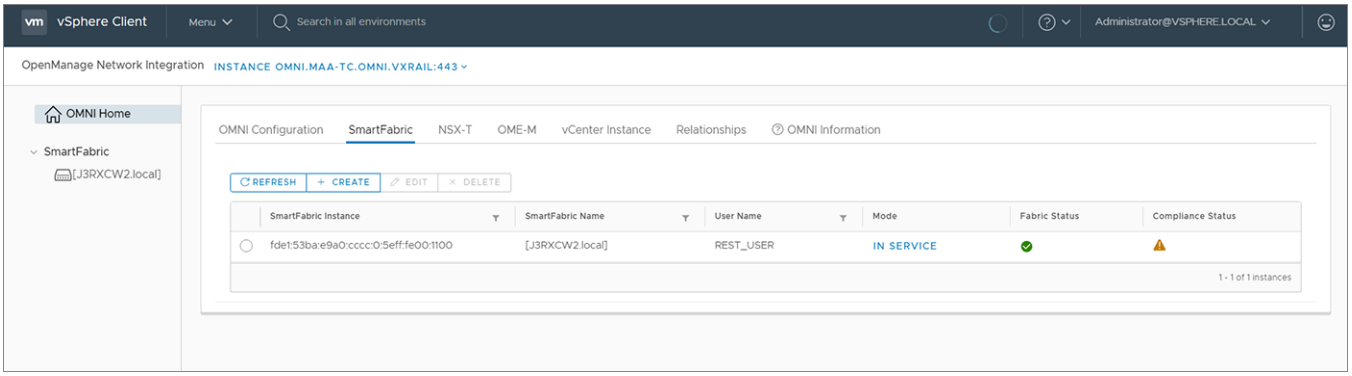

NOTE: During VxRail initial deployment, the system forces you to change the password. If you forget the REST_USER password, contact Dell support to reset REST_USER password.

3. Edit the SmartFabric name, password, or enable or disable SFS events, and click **Submit**.

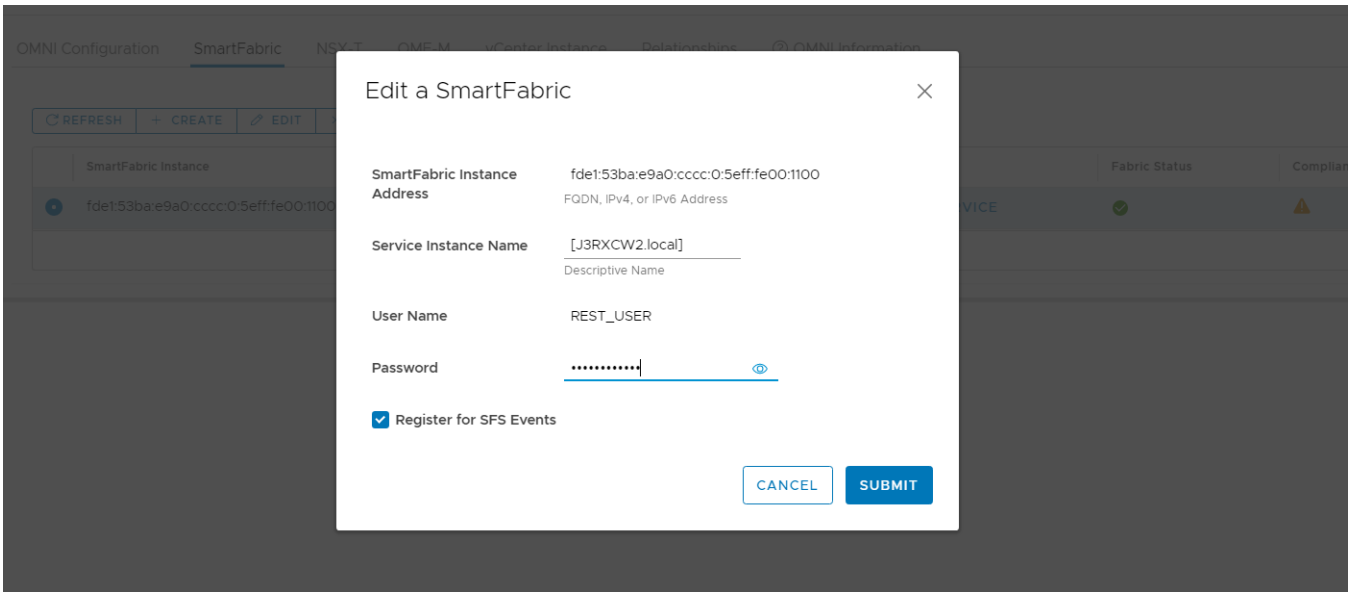

NOTE: SFS events feature is supported from SmartFabric OS10.5.2.2 version or later.

4. The system displays SmartFabric instance configuration success message.

Add SmartFabric instance

This information describes how to add SmartFabric instances in OMNI. You can add up to 15 SmartFabric instance in a single OMNI VM.

1. Identify the master IP address of the switch in a SmartFabric cluster. To identify the master, use the `show smartfabric cluster` command in the OS10 switch CLI.

```
OS10# show smartfabric cluster

-------------------------------------------------------
CLUSTER DOMAIN ID : 100
VIP               : fde2:53ba:e9a0:cccc:0:5eff:fe00:1100
ROLE              : MASTER
SERVICE-TAG       : FX6HXC2
```

```
MASTER-IPV4        : 10.11.180.8
PREFERRED-MASTER   : true
-------------------------------------------------------------
```

2. Go to the OMNI portal.
3. From **OMNI Home** > **SmartFabric**, click **Create** to manually add the master IP address of the SmartFabric instance.

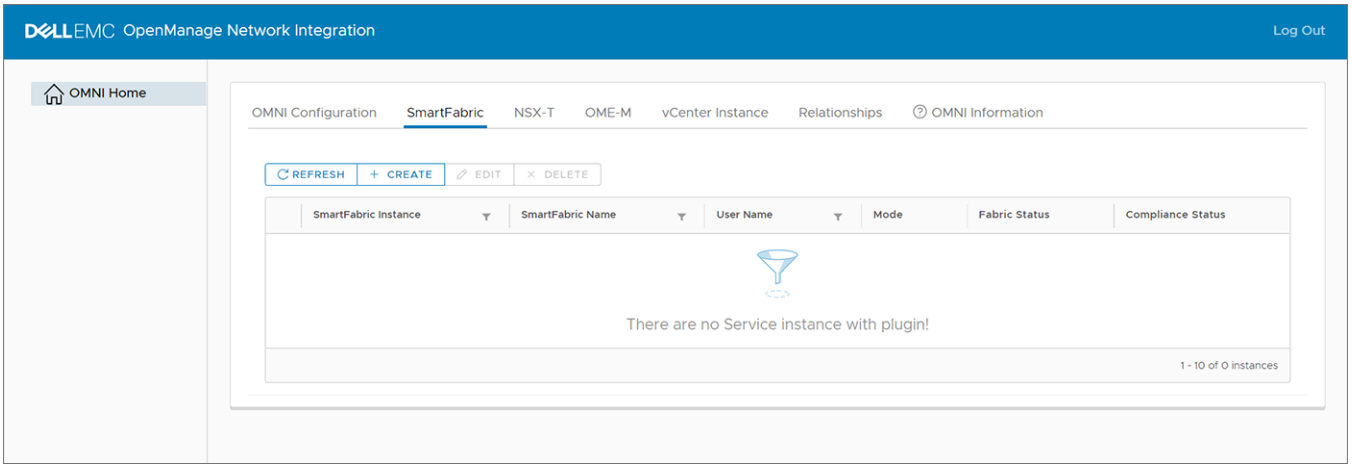

4. Enter the SmartFabric instance IP address, SmartFabric name, username, and password.
5. (Optional) Select **Register for SFS events** checkbox to retrieve the SFS events and display through OMNI. Click **Add**.

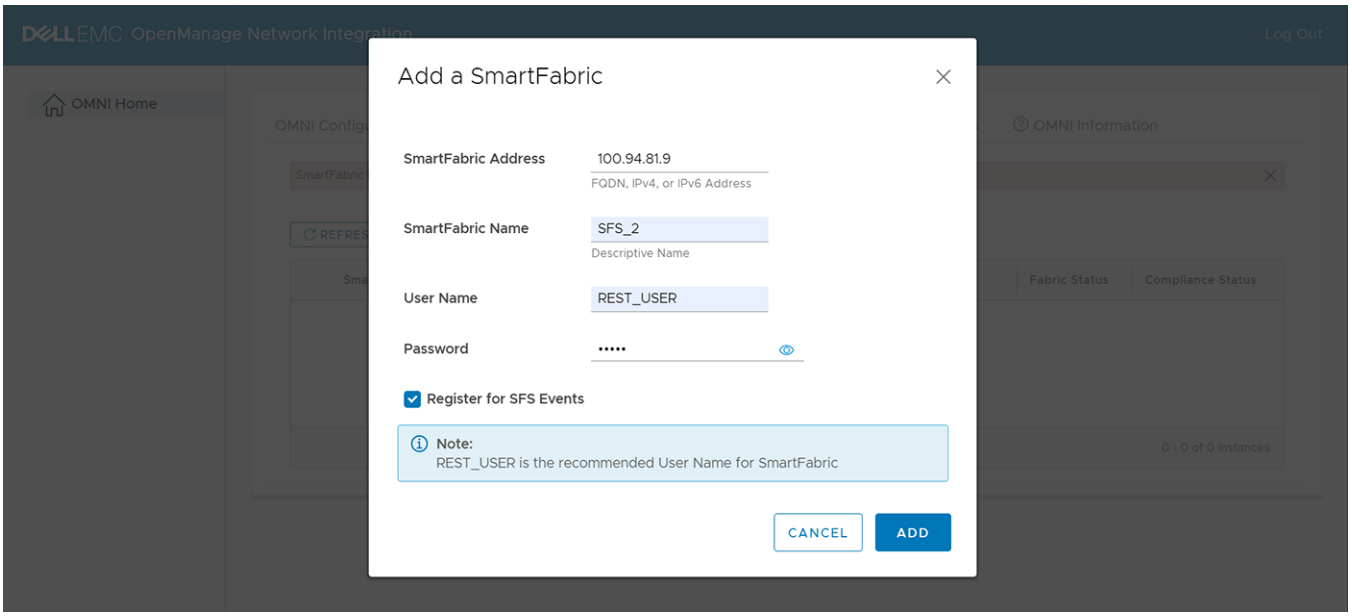

6. The system displays SmartFabric instance creation success message.

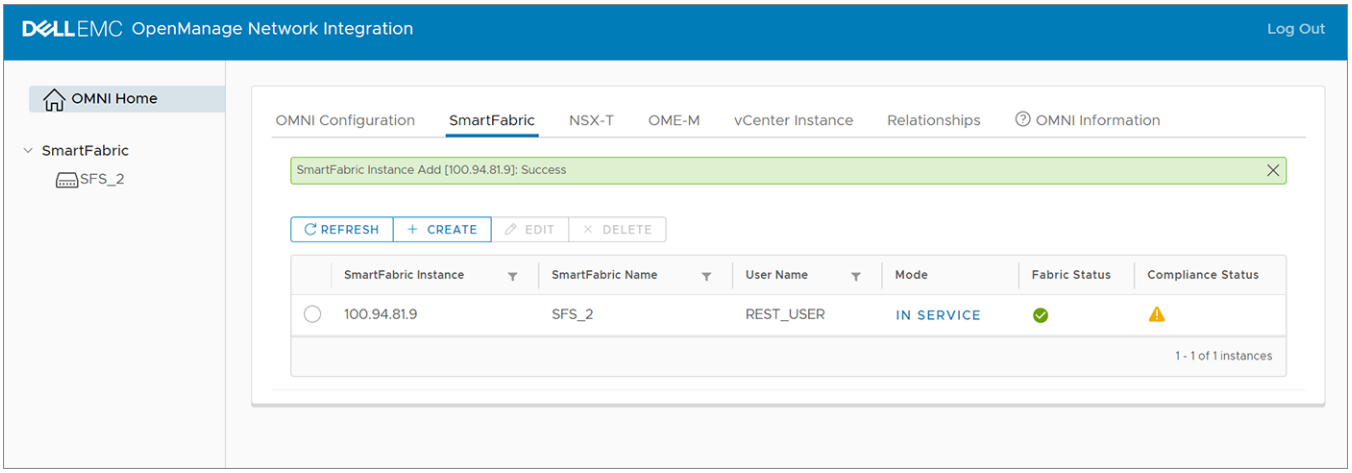

NOTE: The compliance status feature is supported from SmartFabric OS10.5.2.2 version onwards. OMNI displays the compliance status information for the SmartFabric instance only if the version running on the switches is OS10.5.2.2 or later. If not, the compliance status is displayed as N/A.

Edit a SmartFabric

To edit the configuration of the existing SmartFabric instance:

1. Select the SmartFabric instance from the list, and click **Edit**.

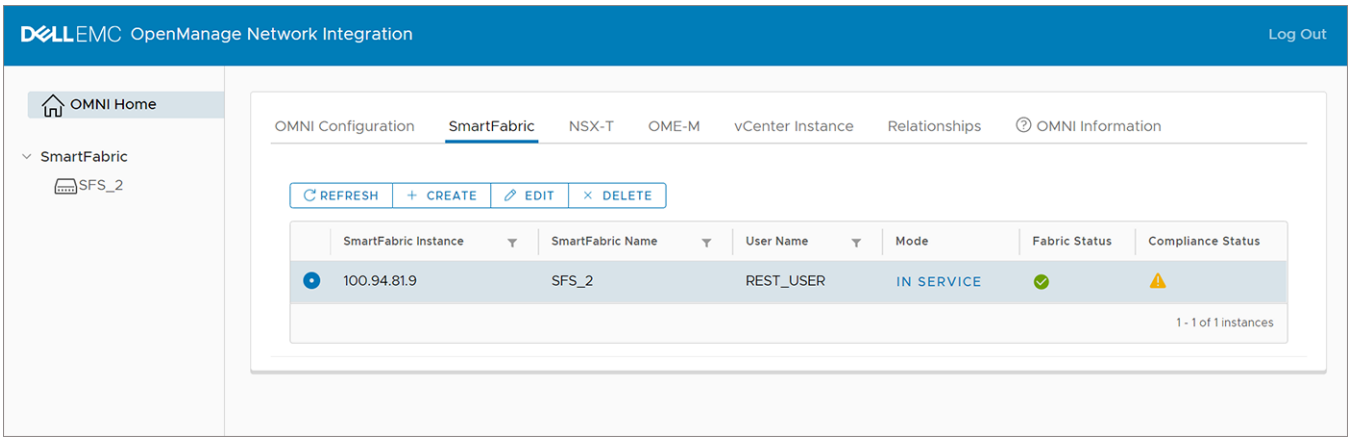

2. Update the configurations and click **Submit**.

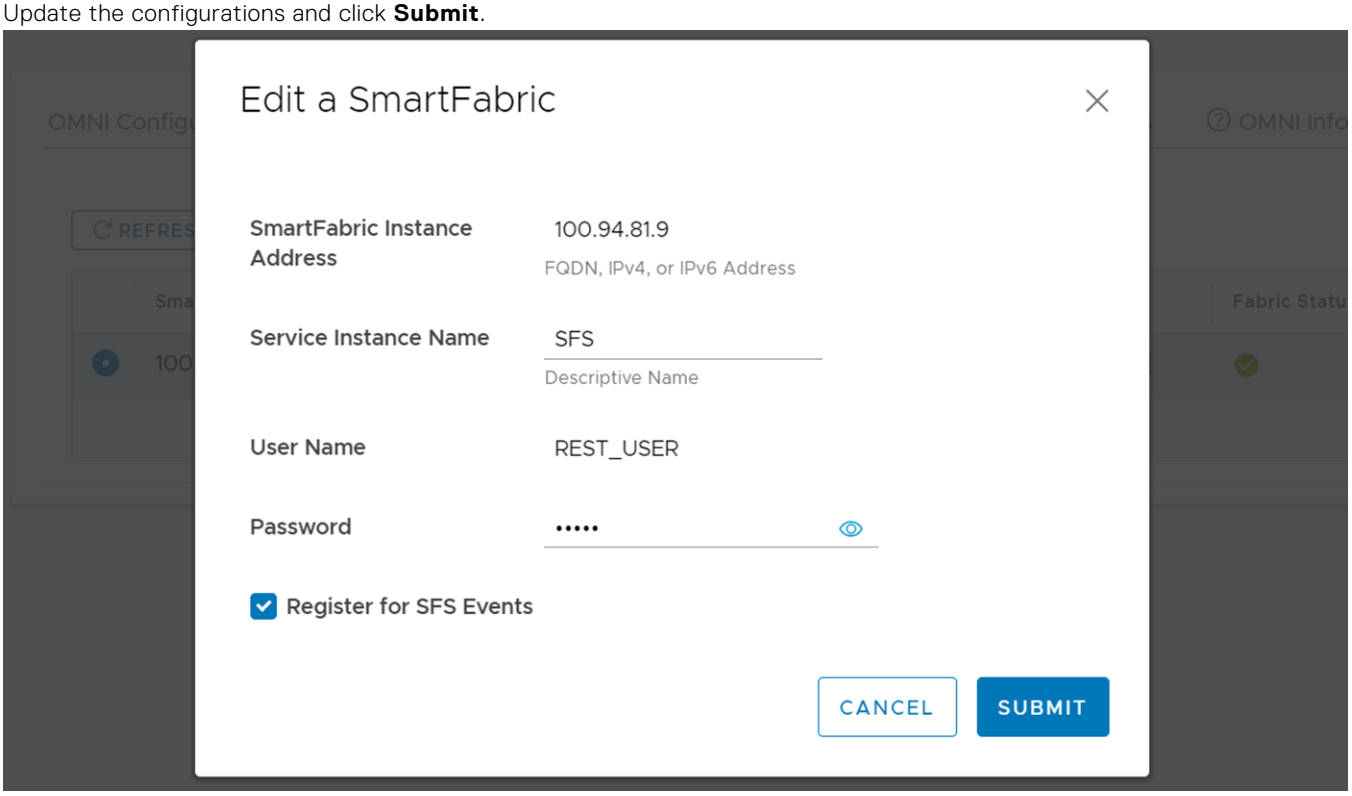

3. The system displays SmartFabric instance update success message.

Delete a SmartFabric

To remove a SmartFabric instance from OMNI.

1. Select the SmartFabric instance from the list and click **Delete**.

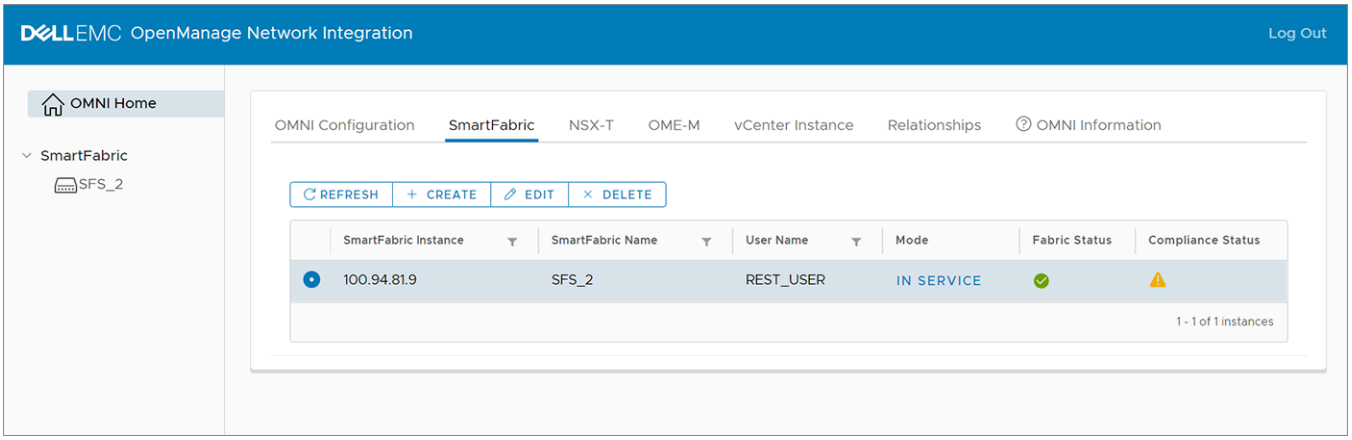

2. Click **Delete** to confirm.

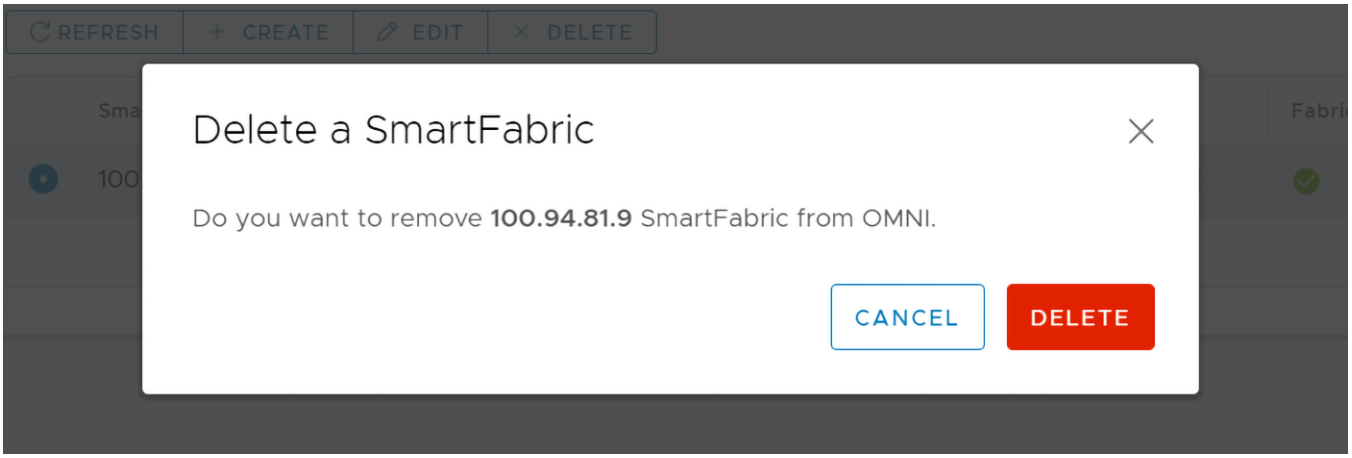

3. The system displays SmartFabric instance update success message.

SmartFabric Maintenance mode

Enabling Maintenance mode prevents OMNI from configuring networks on SmartFabric when there are changes in the vCenter port groups and disables the UI navigation for that instance.

The **OMNI Home** > **SmartFabric** page displays the mode of each SmartFabric instance added in the OMNI VM.

Enable Maintenance mode

To enable Maintenance mode for a SmartFabric instance:

1. From **OMNI Home** > **SmartFabric**, click **In Service** mode for a specific SmartFabric instance.

2. Click **Ok** to confirm.

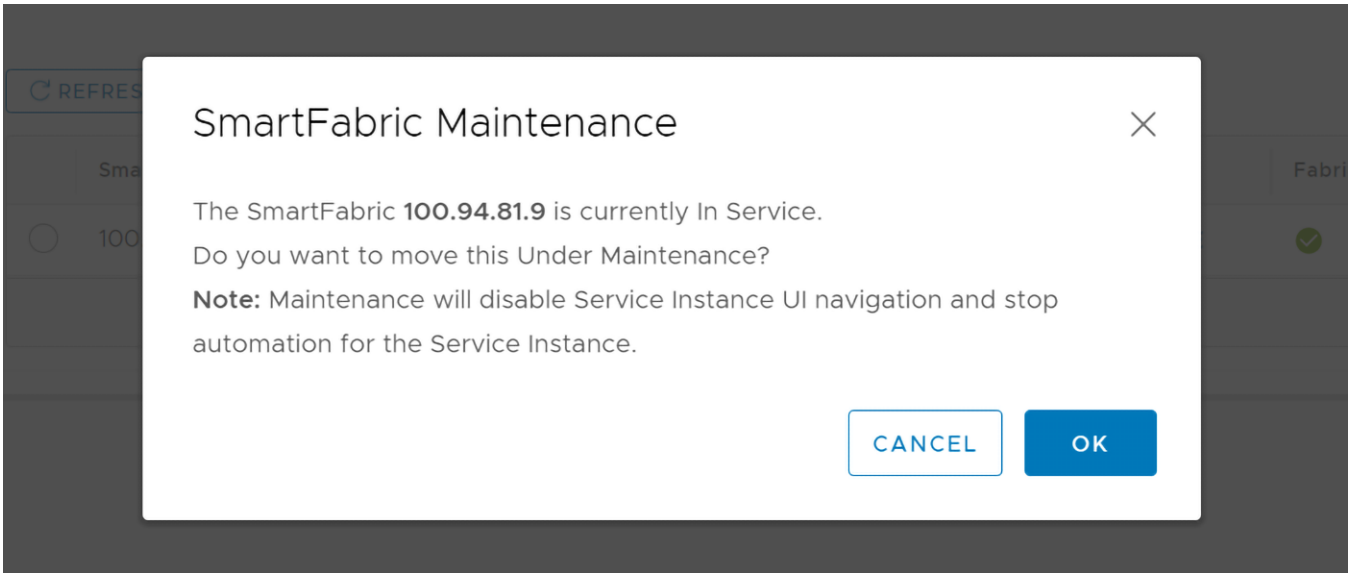

The SmartFabric instance mode changes to **Under Maintenance**. During Maintenance, the SmartFabric UI navigation is disabled and the automation services is stopped for the instance.

Disable Maintenance mode

To disable Maintenance Mode for a SmartFabric instance:

1. From **OMNI Home** > **SmartFabric**, click **Under Maintenance** for a specific SmartFabric instance.
2. Click **Ok** to confirm. The SmartFabric instance mode changes to **In Service**.

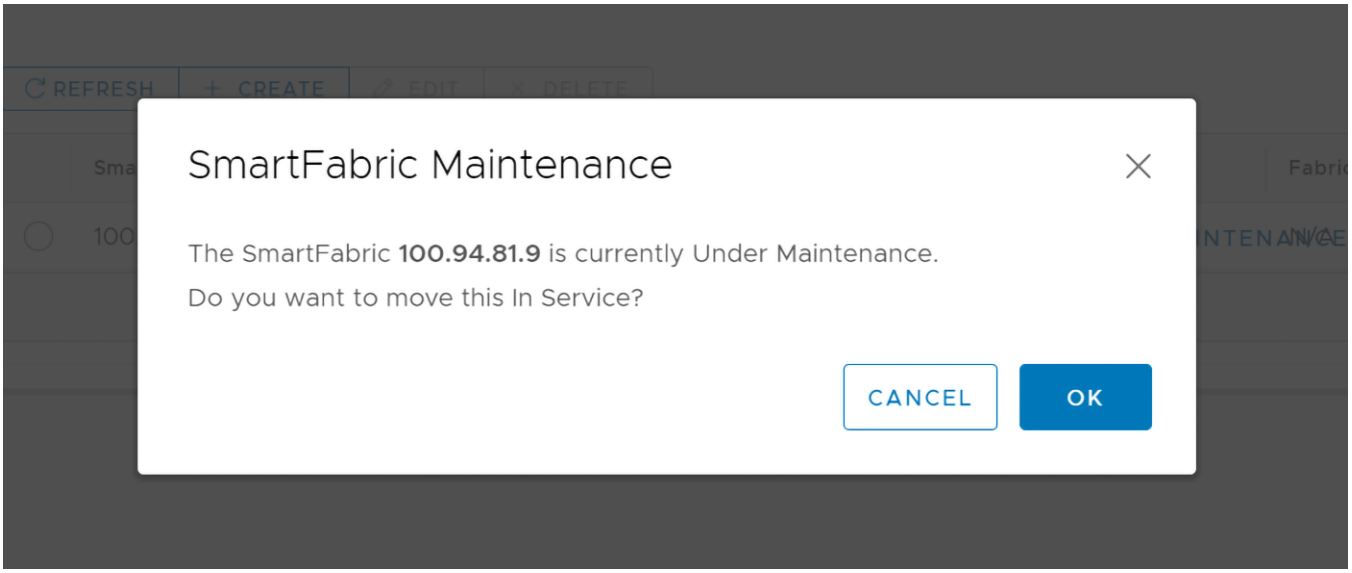

OMNI support for vCenter Enhanced Linked mode

Enhanced Linked mode (ELM) is a feature available in vCenter. Using ELM, you can link multiple vCenter appliances that are deployed across different location and have a global view of the inventory.

OMNI appliance behavior when the vCenter or vCenters registered to OMNI are in ELM:

- You must register all the vCenters that are in ELM with OMNI. For example, if two vCenters vCenter1 and vCenter2 are linked using ELM, you must register both the vCenters (vCenter1 and vCenter2) to launch OMNI plug-in from vCenter1 and vCenter2. For more information, see Register OMNI with vCenter.
- If you want to launch OMNI plug-in from a vCenter that is in ELM and does not have any host that is connected to SmartFabric instance, you can only register the vCenter and disable automation. To do that, select True for registration option and False for automation option when adding the vCenter instance in OMNI. For more information, see Register OMNI

with vCenter. In this example, if vCenter2 does not have any host that is connected to SmartFabric instance added to OMNI, you can only register the vCenter and disable the automation.

- Dell Technologies recommends using stand-alone OMNI UI to unregister all the vCenters that are linked through ELM.

Host network inventory

You can view information about physical Dell EMC PowerSwitch infrastructure running SmartFabric OS10.

Host network inventory page

Select a host in vCenter, select the **Monitor** tab, then select **OpenManage Network Integration** (OMNI) in the monitor sidebar.

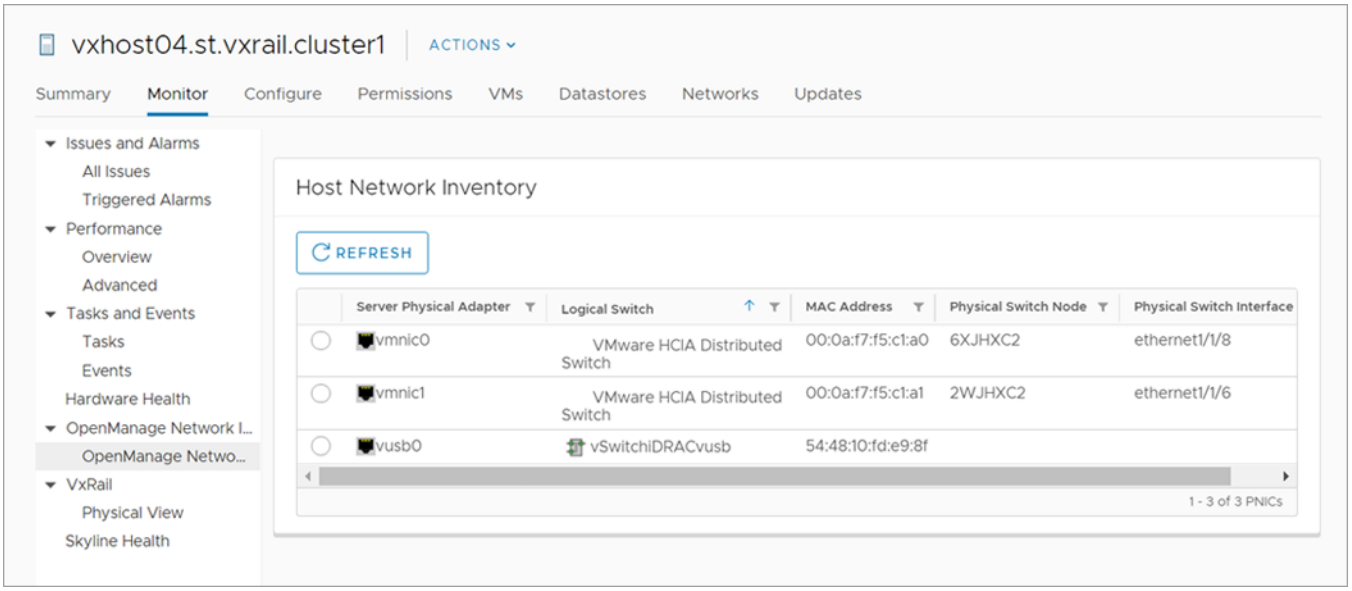

Refresh button

Click **Refresh** to update the host network inventory data and display updated contents.

Physical adapter table

Select a switch from the Host Network Inventory to view detailed information. The table is default-sorted by descending switch name to group physical adapters belonging to the same switch.

- Server Physical adapter—Name of the physical NIC.
- Logical switch—Name of switch the physical adapter is connected to.
- MAC address—AC address of the physical adapter.
- Physical switch node—Service tag of physical switch that is connected to the fabric.
- Physical switch interface—Physical switch port this server NIC is connected to.

View logical switch details

Displays information about the logical switch that is connected to the selected physical adapter.

When you select a server physical adapter from the Host Network Inventory, the page displays the information about logical switch that is connected to the selected physical NIC.

- Switch tab—Includes name of switch, MTU in bytes of switch, physical adapters connected to the switch, and uplink ports on the switch.

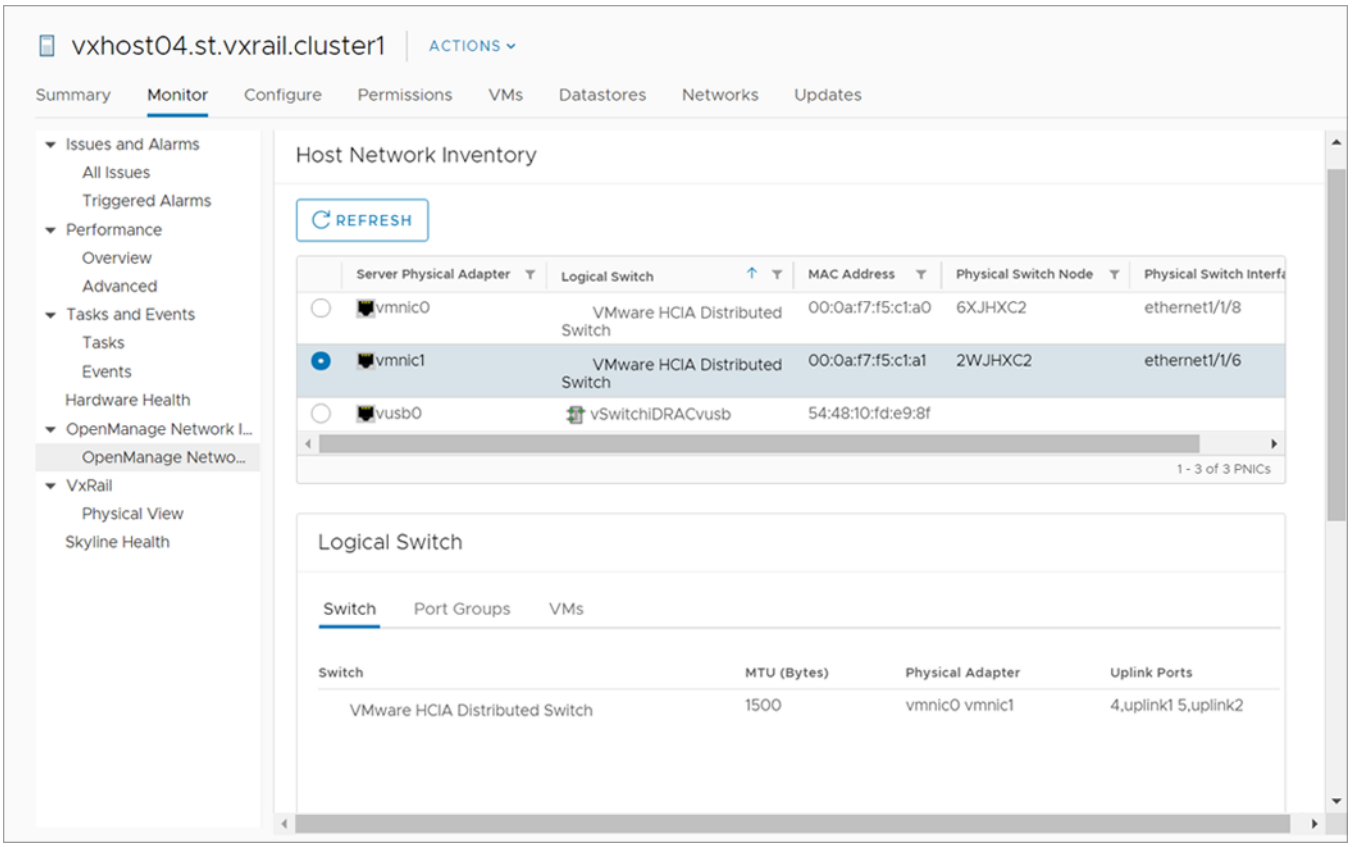

- Port groups tab—Includes the name of port groups, and VLAN IDs for each port group.

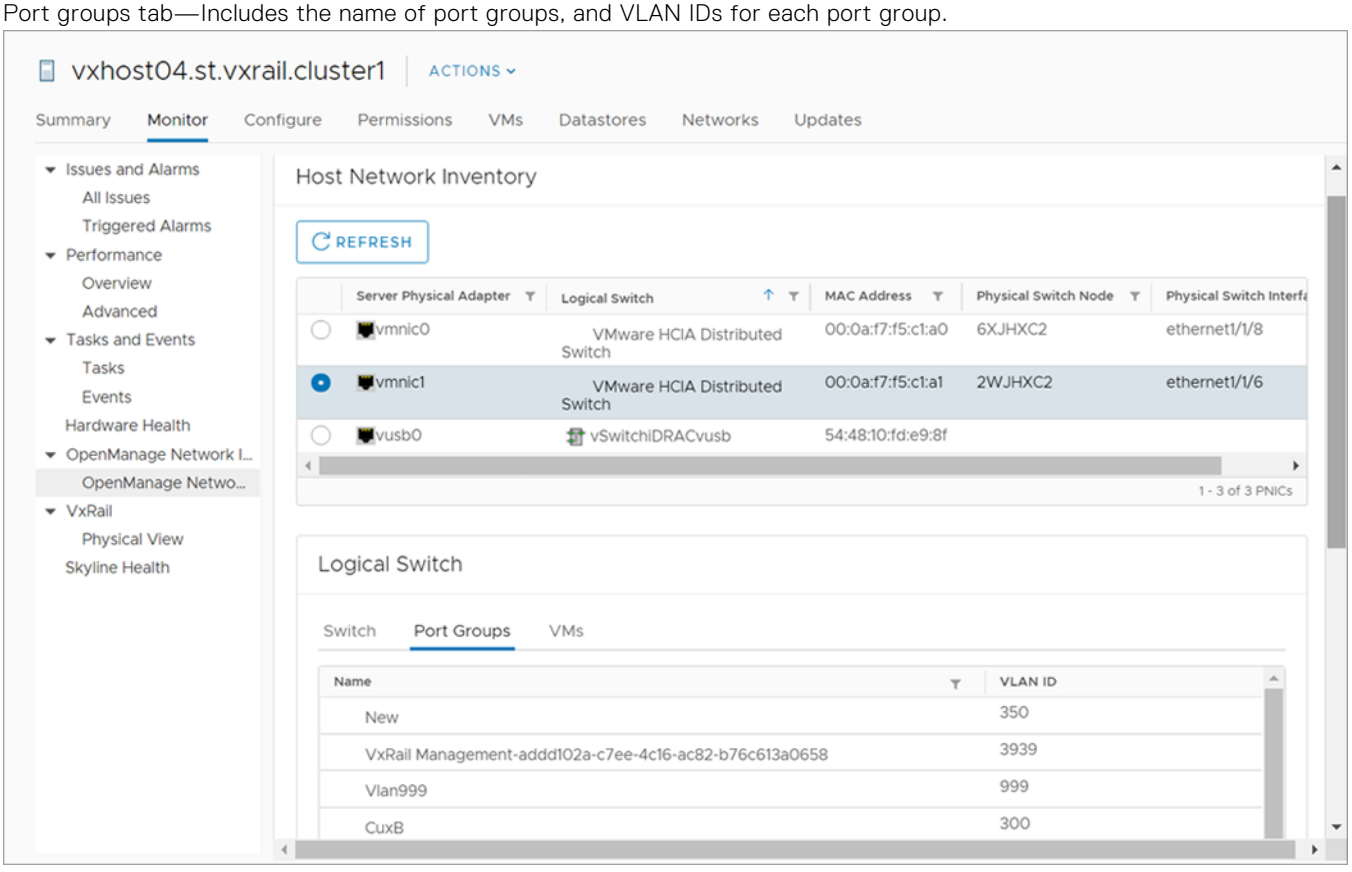

- VMs tab—Includes the name of VMs of that host that is connected to a single virtual switch.

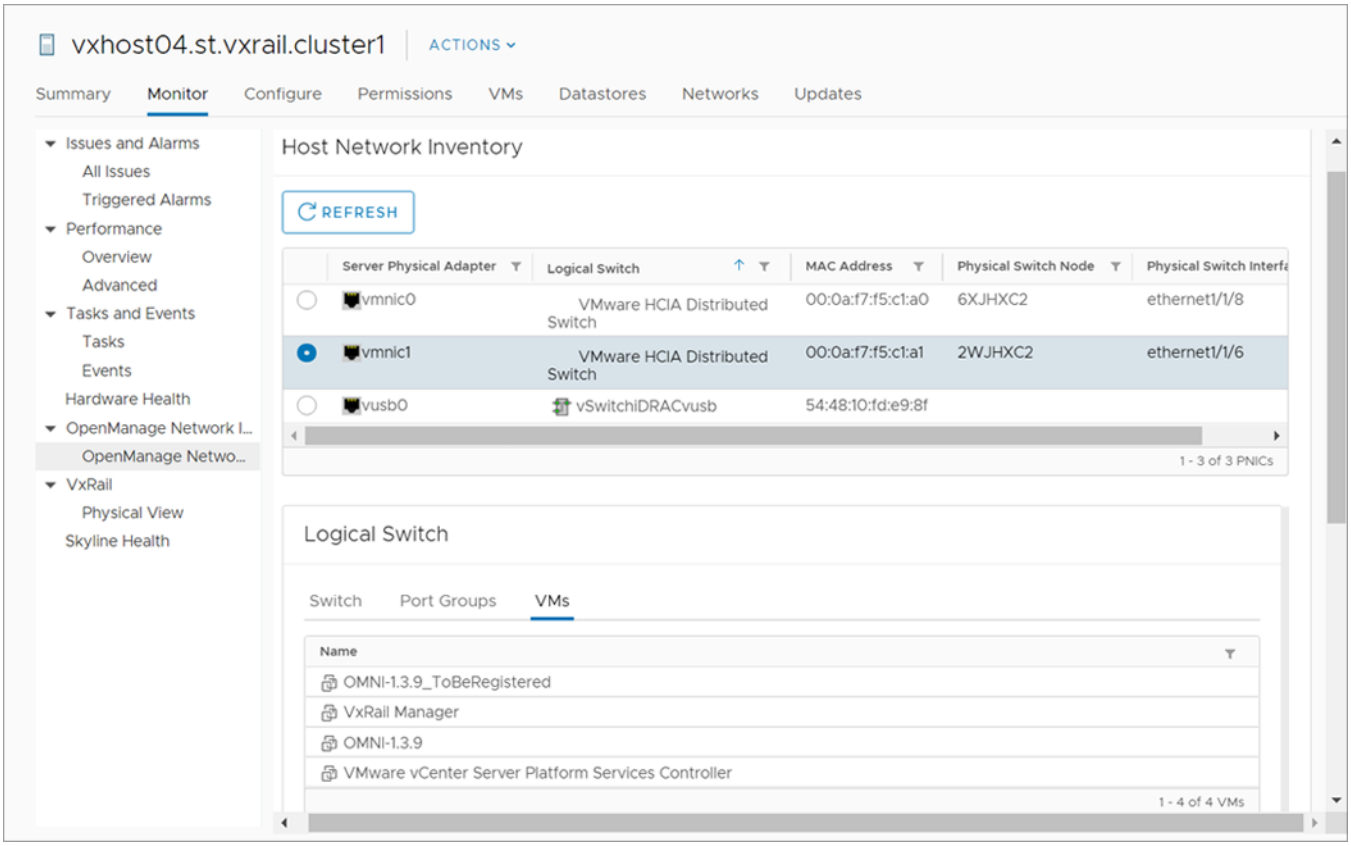

View physical switch details

Displays information about the onboard interface. This information displays only when there is a physical connection between the VxRail domains and OMNI.

When you select a server physical adapter from the Host Network Inventory, the page also displays the information that is related to physical switch connected to the selected physical NIC.

Onboard interface tab

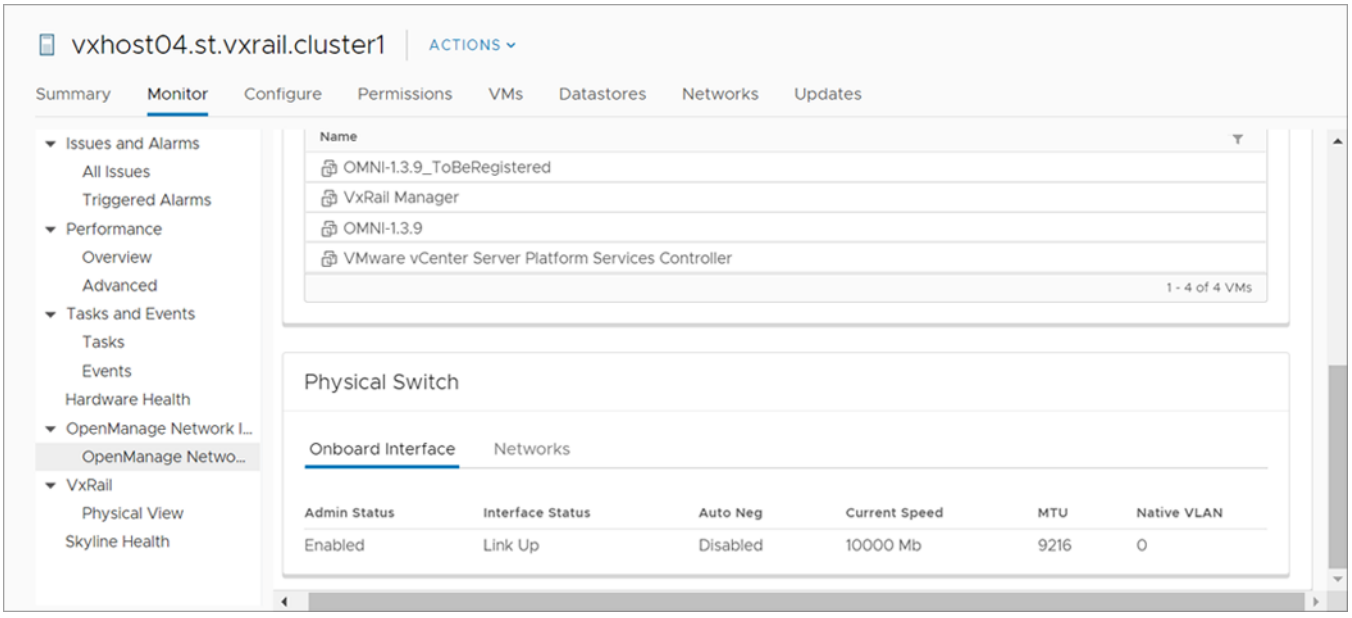

- Admin Status—configured state of the physical interface
- Interface Status—current operations state of the physical switch port

- Auto Neg—negotiation status of the physical interface
- Current Speed—current operational speed of the physical interface
- MTU—maximum transmitting unit configured on the physical interface
- Native VLAN—untagged default VLAN for the physical switch

Networks tab

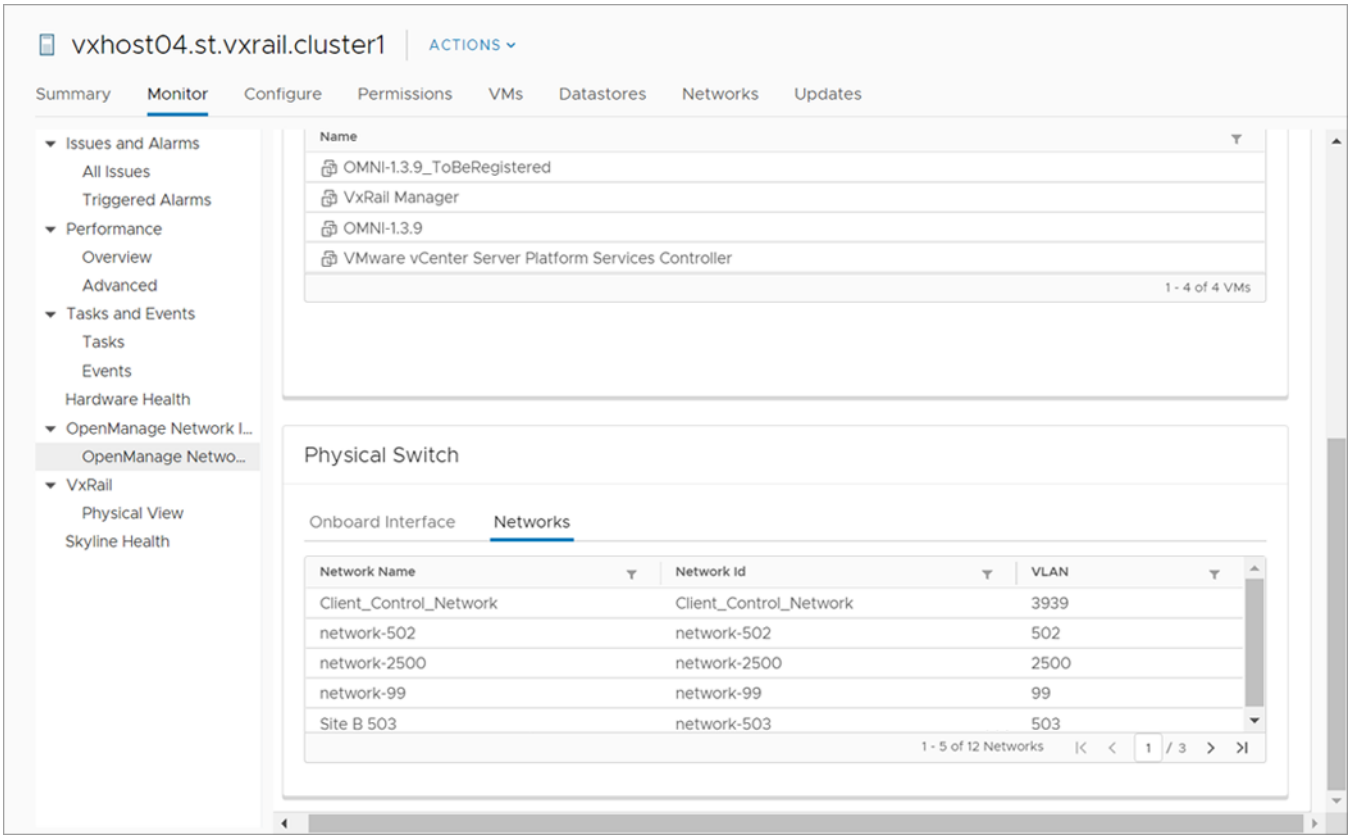

- Network Name—name of the VLAN network
- Network ID—unique identifier of the fabric network
- VLAN—tagged VLAN of the switch port

View service instance and vCenter relationships

Starting from 2.0 release, OMNI displays the relationships between the vCenter and the service instances (SFS, OME-Modular, or NSX-T), and the relationship type.

OMNI automation service periodically queries the hosts information from the service instances and the registered vCenters. The information is used to build the relationship between the service instances and the vCenter. Select **OMNI Home** > **Relationships** tab to view the relationship information and the type of the relationship between the entities.

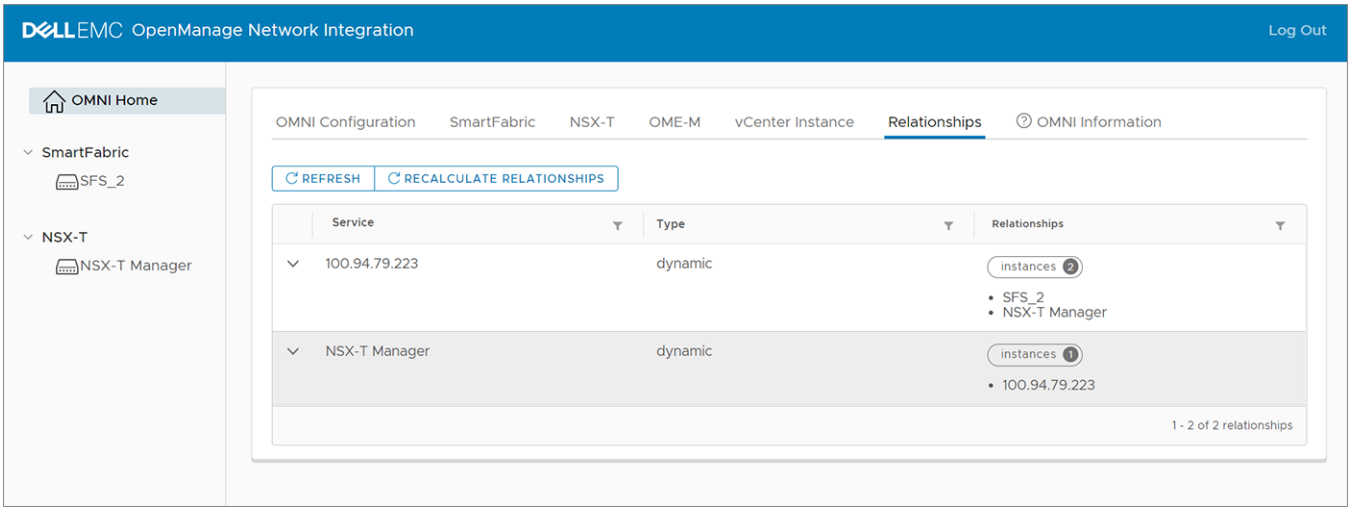

OMNI Information

You can view the links to documentation and end-user license agreement (EULA).

Plugin Information Links has links to:

- Documentation
- EULA

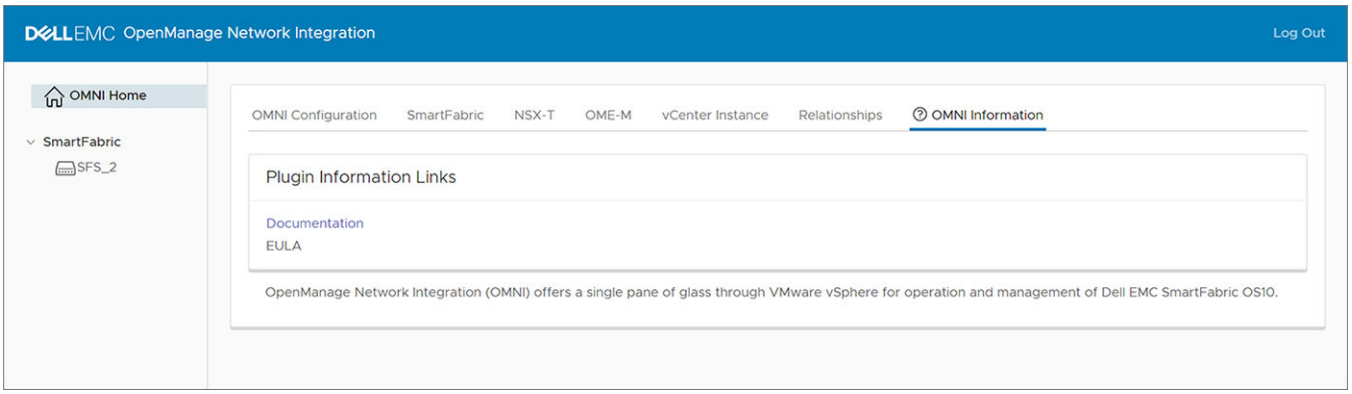

1. Click **EULA** to view the end-user license agreement.

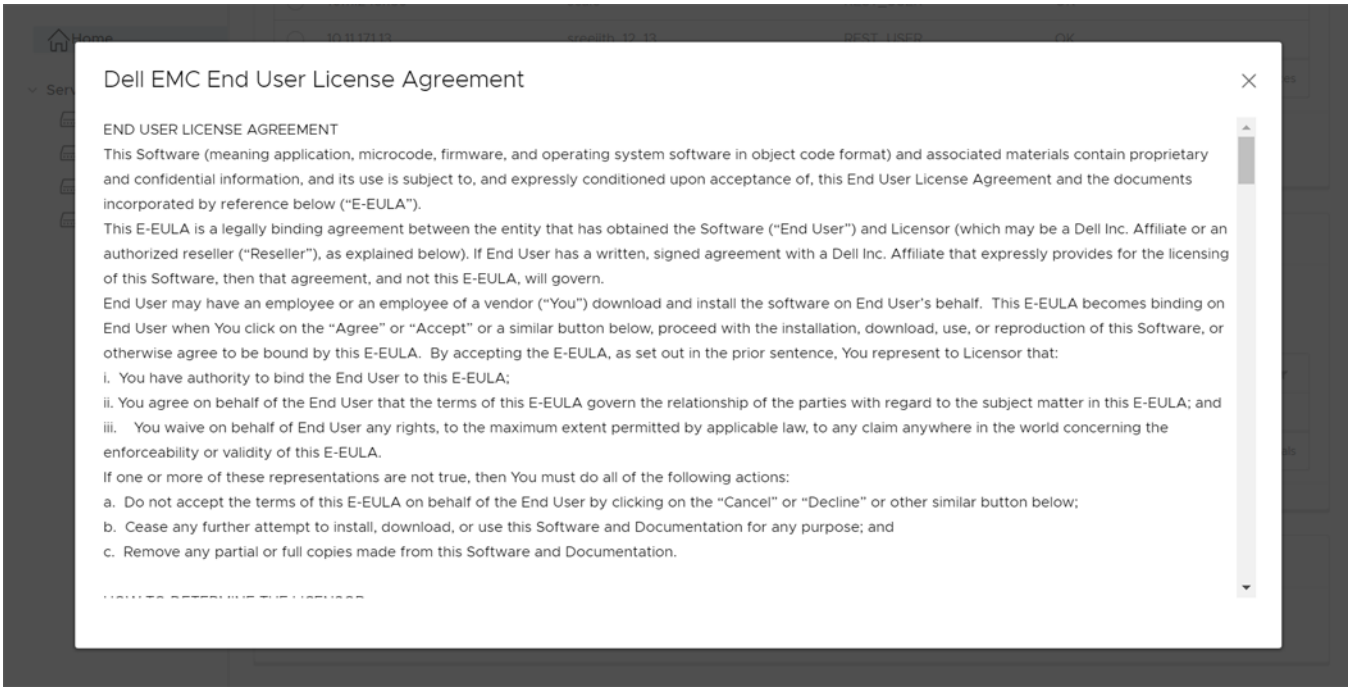

2. Click **Documentation** to see the documents that are uploaded at www.dell.com/support OpenManage Network Integration product page.

OMNI Appliance Management user interface

From OMNI 1.3 release, a new UI—OMNI Appliance Management is introduced to manage all the system, web, and automation services running in the OMNI.

After you create the OMNI virtual appliance and complete the virtual appliance setup, you can launch the OMNI appliance management UI.

You can access the OMNI Appliance Management UI from the OMNI stand-alone page, see Access OMNI stand-alone portal. Click **Launch OMNI Appliance Management** link from the page.

NOTE: Access OMNI Appliance Management UI only with OMNI VM appliance administrator credentials.

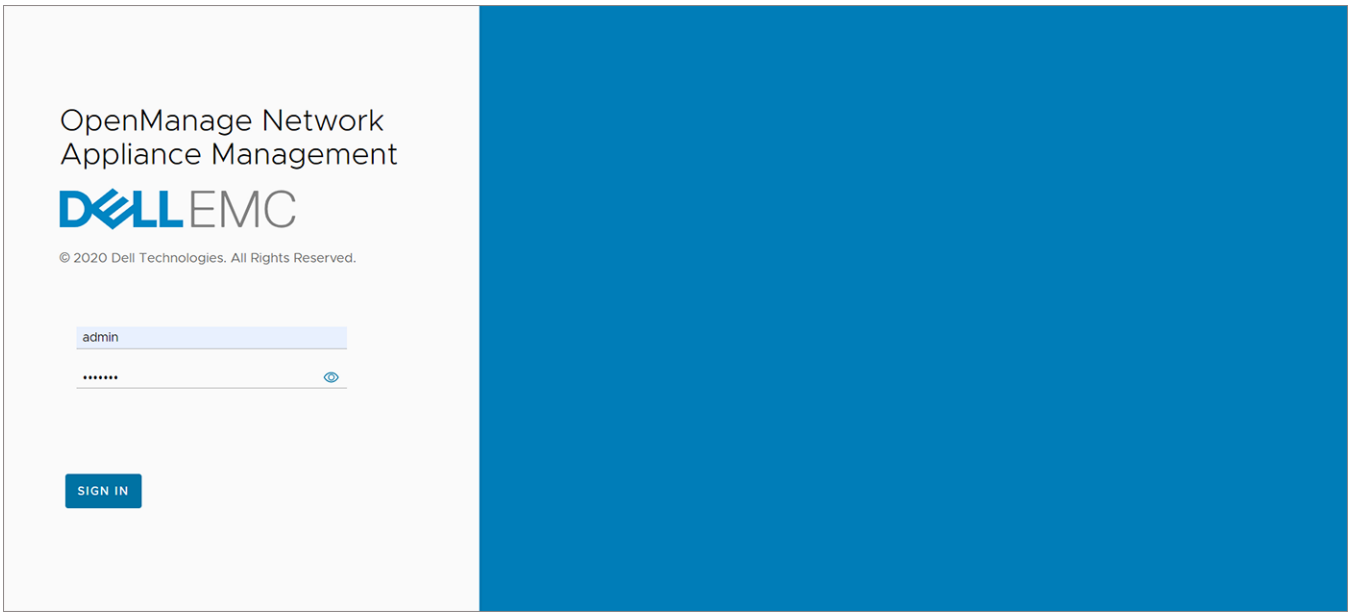

NOTE: You can also access the Appliance Management UI directly from a browser. Open a browser, go to **https://*OMNI_IP*/omni** with the IP address or FQDN that is configured during the initial setup.

- **Logout**—Manually terminate the login session using the **Log out** button in the upper right of the UI.
- **Login session timeout**—OMNI terminates an inactive login session after 15 minutes to prevent unauthorized access.

View OMNI Appliance Management summary

Summary page displays:

- Appliance Version—Displays the build number of OMNI VM appliance (OMNI OVA version deployed).
- Hostname—Displays the hostname configure during OMNI setup.
- Product—Displays the name of the VM appliance that is registered with the vCenter.
- Software Version—Displays the version of the OMNI VM build.

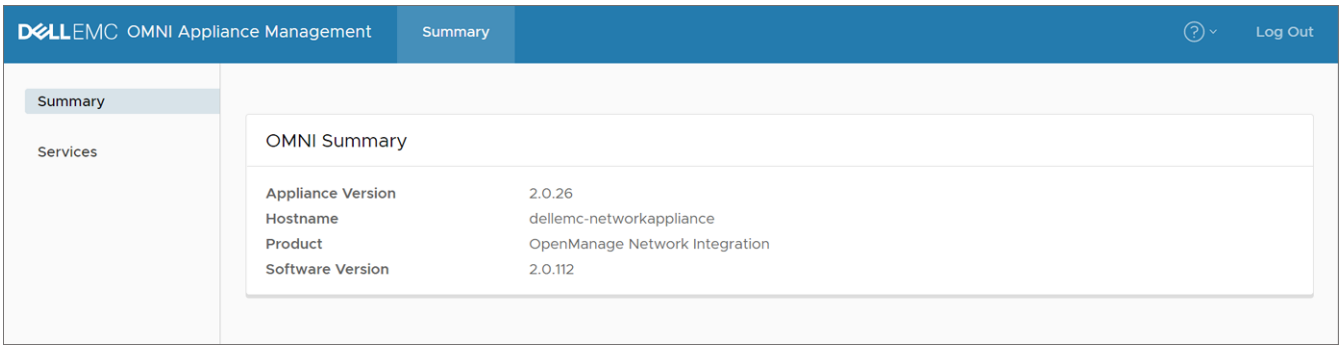

Manage OMNI essential and automation services

Services menu displays all the management and vCenter automation services running on the OMNI appliance.

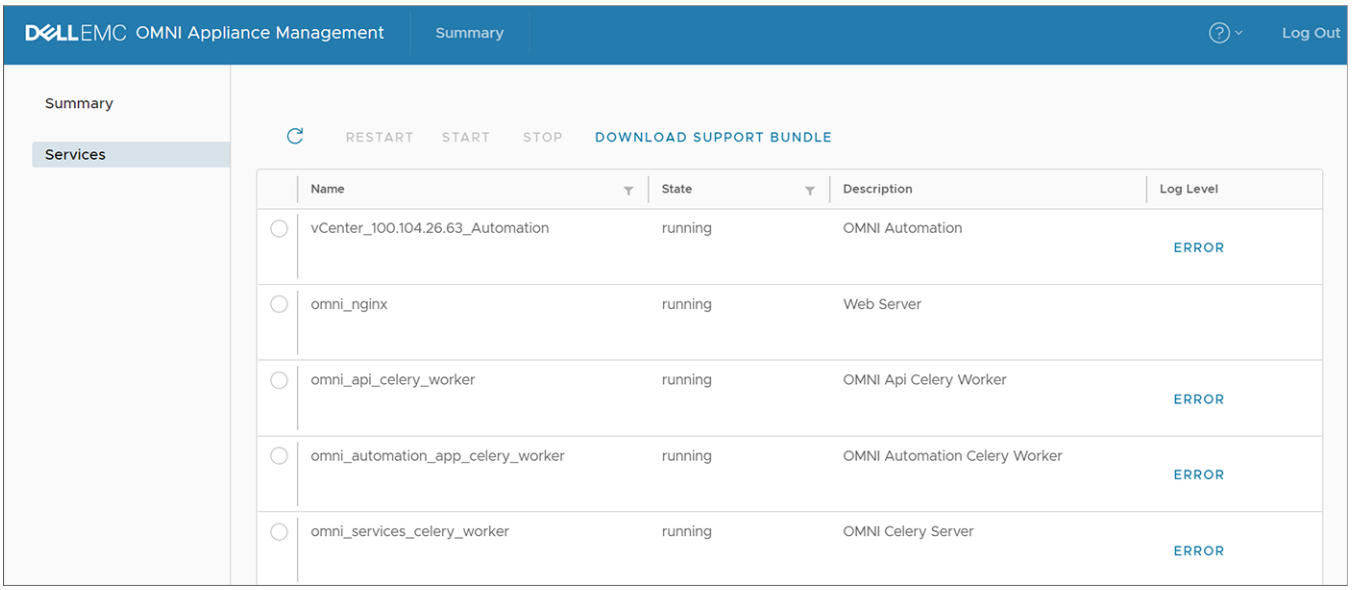

By default, the web and database essential services start automatically after the initial setup. After adding the SmartFabric, OME-M, or NSX-T instances and registering the relevant vCenters, OMNI creates automation services for each vCenter instance. Automation services that are related to the SmartFabric , OME-M, or NSX-T instances start depending on the automation option set during the registration of the vCenter.

Table 10. List of OMNI services

Service	Function	States
omni_api	Service serving REST APIs for OMNI Fabric Management interface.	Can restart the services.
omni_services	Orchestration service that provides APIs to start, stop, and manage all OMNI services.	
omni_events_receiver	Events receiver service receives events from the SFS and store in the message queue.	
omni_api_celery_worker	Worker service that conducts fabric upgrades and vCenter re-registration when registration data is updated.	
omni_automation_app_celery_worker	Automation task service that identifies vCenter configuration change tasks and synchronizes all hosts that have been changed on the vCenter.	
omni_services_celery_worker	The OMNI services celery worker manages automation container startup after OMNI services are started or restarted.	
omni_events_celery_beat	Service that periodically cleans the old events from the database.	
omni_events_celery_worker	Worker service that process the events from the message queue and stores them in the database.	
omni_automation_app_celery_beat	Service that periodically prunes unused networks on service instances and discovers how service instances are related to each other.	
omni_queue	Service that runs the message queue. This service enables communication between other services and also to add and perform celery tasks.	
omni_db	Database service that stores crucial information.	Cannot restart, start, or stop the service.
omni_nginx	Web server service that manages all incoming and outgoing web requests.	
vCenter automation services	Automation services running for each vCenter	Can start, stop, or restart the services.

Click **Refresh** icon to update the data and display the updated contents.

Start vCenter automation services

To start the fabric automation services:

1. From the **OMNI Appliance Management** UI, click **Services** tab menu.
2. Select the automation service that you want to start, and click **Start**.

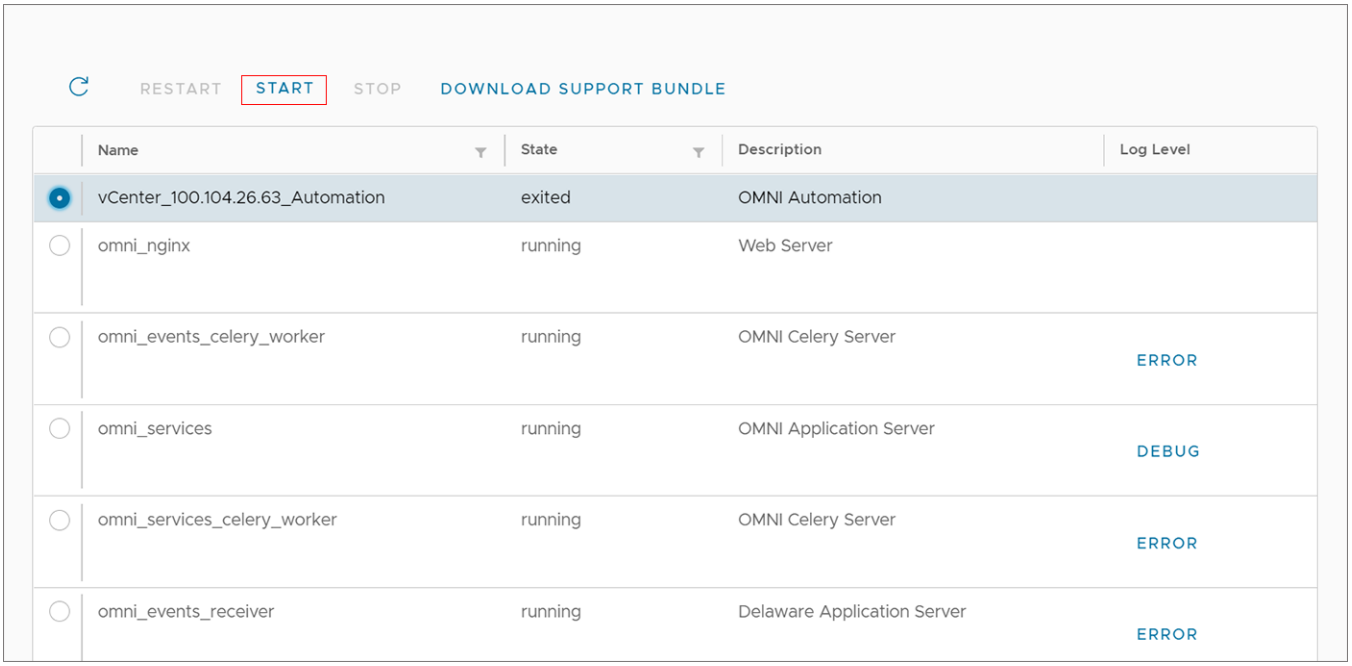

After you start the service, OMNI starts monitoring the networking events for the registered vCenter.

3. The system displays start service success message.

Stop vCenter automation services

To stop the fabric automation services:

1. Select the relevant automation service that you want to stop, and click **Stop**.

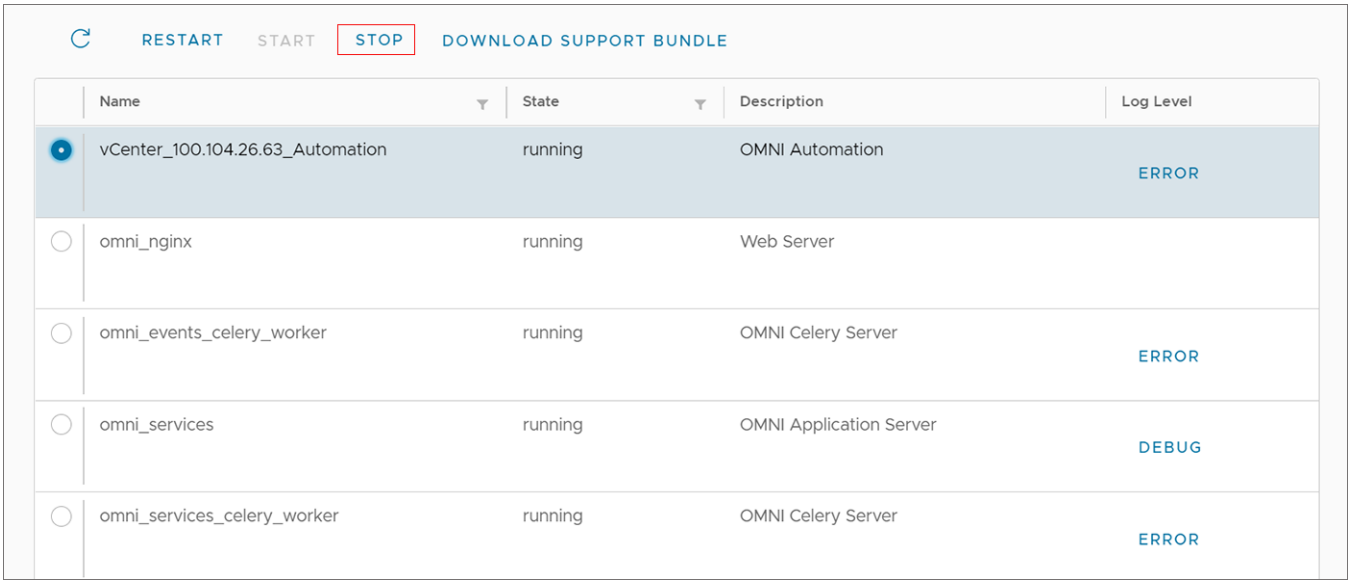

2. The system displays stop service success message.

To restart the fabric automation service, select the relevant automation service, and click **Restart**.

Change log level

1. When the log-level of OMNI is set to ERROR, the system records the error logs. When the log-level is set to DEBUG, error logs and logs with additional information is recorded. Use the DEBUG level when you want to diagnose an issue.

2. (Optional) Click **Error** under log-level of each service to modify the log-level to **Debug**.

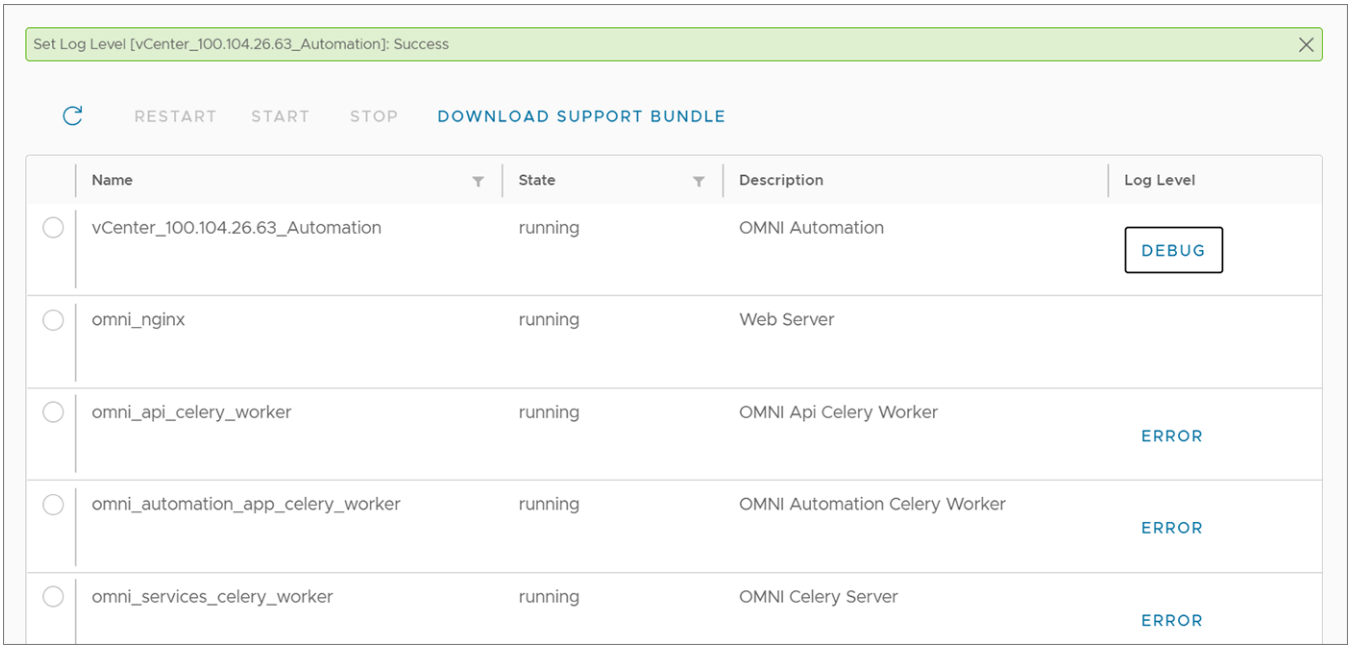

The system displays set log level success message.

3. (Optional) Click **Debug** under log-level of each service to modify the log-level to **Error**.

   The system displays set log level success message.

Download Support Bundle

Support options are used for debugging. If there is an issue, download a support bundle containing all the logs that are found in OMNI. Also change the log-level of OMNI to collect logs of different types.

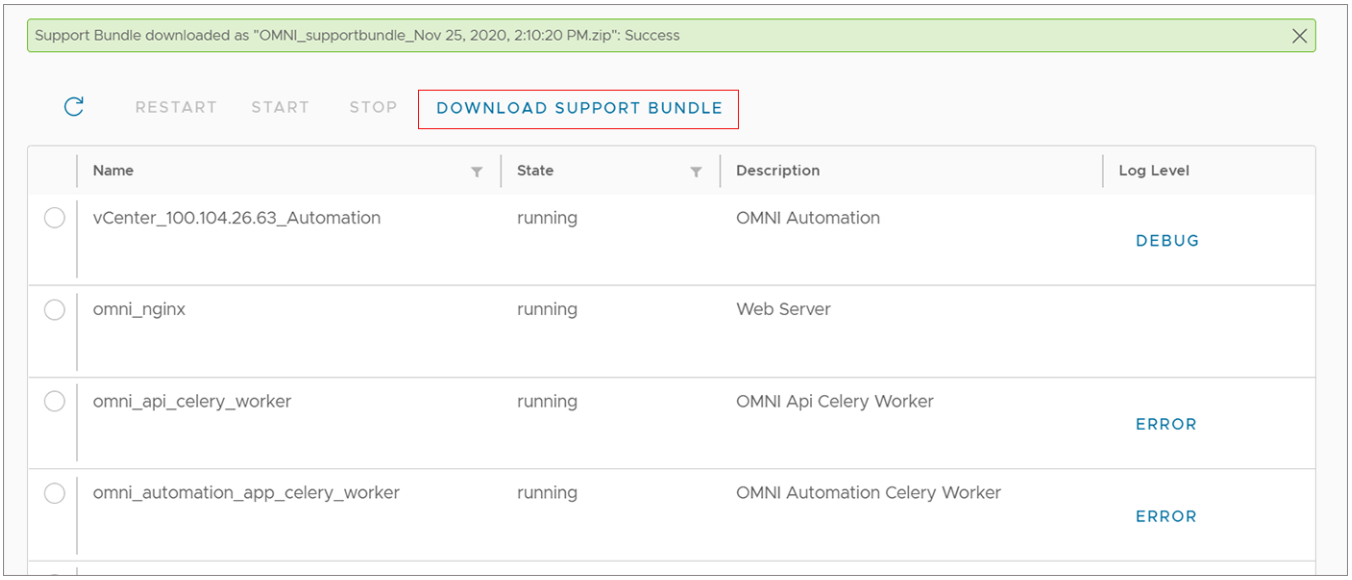

Help links

Using the help icon, you can:

- Access the Dell EMC OpenManage Network Integration documentation support page.
- View the end-user license agreement (EULA).

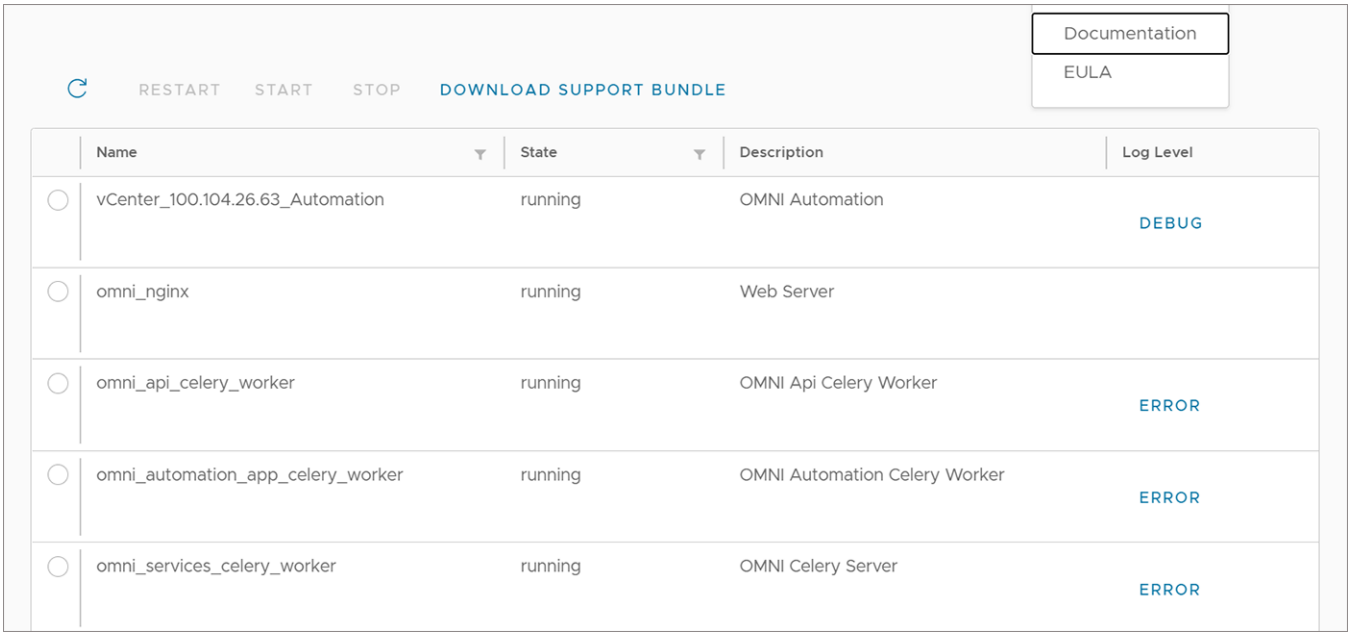
Documentation
EULA
RESTART
START
STOP
DOWNLOAD SUPPORT BUNDLE
Name
State
Description
Log Level
vCenter_100.104.26.63_Automation
running
OMNI Automation
DEBUG
omni_nginx
running
Web Server
omni_api_celery_worker
running
OMNI Api Celery Worker
ERROR
omni_automation_app_celery_worker
running
OMNI Automation Celery Worker
ERROR
omni_services_celery_worker
running
OMNI Celery Server
ERROR

7

OMNI automation support for PowerEdge MX SmartFabric

Starting from 2.0 release, OMNI manages fabric automation for ESXi hosts deployed within the Dell EMC PowerEdge MX solution when running SmartFabric Services. For any change to the port group configuration in vCenter, OMNI automatically associates the VLAN to the applicable host-connected ports on the switch. OpenManage Enterprise Modular (OME-Modular) is embedded in the MX platform and enables configuration and management of up to 20 MX7000 chassis from one interface. For more information about SmartFabric services on PowerEdge MX, see PowerEdge MX documents.

View the logical and physical switch inventory of MX servers in vCenter **Host Network Inventory** page. For more information, see View host inventory.

Prerequisites

Ensure that the following prerequisites are met to support OMNI automation services for PowerEdge MX:

- MX system is healthy with no failed components.
- The PowerEdge MX network switches must be configured in SmartFabric mode and operational.
- The entire PowerEdge MX system must be on the MX 1.20.10 baseline or later.
- MX servers that are considered for automation must be deployed through OME-Modular server profiles.
- The OME-Modular server template should include vCenter infrastructure VLANs such as management and vMotion but not virtual machine VLANs.
- MX servers must have ESXi installed and be connected to the target vCenter.
- Dell Technologies recommends using VMware ESXi version 6.7 and later.
- OMNI must be able to communicate with OME-Modular and vCenter to provide automation.

NOTE: OMNI automation does not support MX servers with NIC partitioning enabled, with the exception of FCoE or iSCSI storage partitions.

For more information about VMware ESXi and PowerEdge MX, see Dell EMC PowerEdge MX VMware ESXi with SmartFabric Services Deployment Guide.

Workflow to integrate OME-Modular with OMNI

Ensure that the prerequisites are met before starting the workflow to integrate OME-Modular with OMNI. Dell Technologies recommends creating a dedicated OME-Modular user account (OMNI_USER) in OME-M for OMNI with a role of Fabric Manager.

NOTE: Do not use the `root` user OME-Modular credentials.

1. Add the OME-Modular service instances in OMNI.
2. Register the vCenters to which the MX servers are connected.
3. Manage automation services for OME-Modular.

Add OME-Modular instance

To manage MX SmartFabric automation using OMNI, add the OME-Modular instance to OMNI. In 2.0 release, you can add one OME-Modular instance in a single OMNI VM.

To add an instance:

1. From **OMNI Home** > **OME-M**, click **Create** to create an OME-Modular service instance by adding the IP address or DNS name of the lead chassis.

   If the OME-Modular instance IP address is a virtual IP address, use the virtual IP address to create the instance.

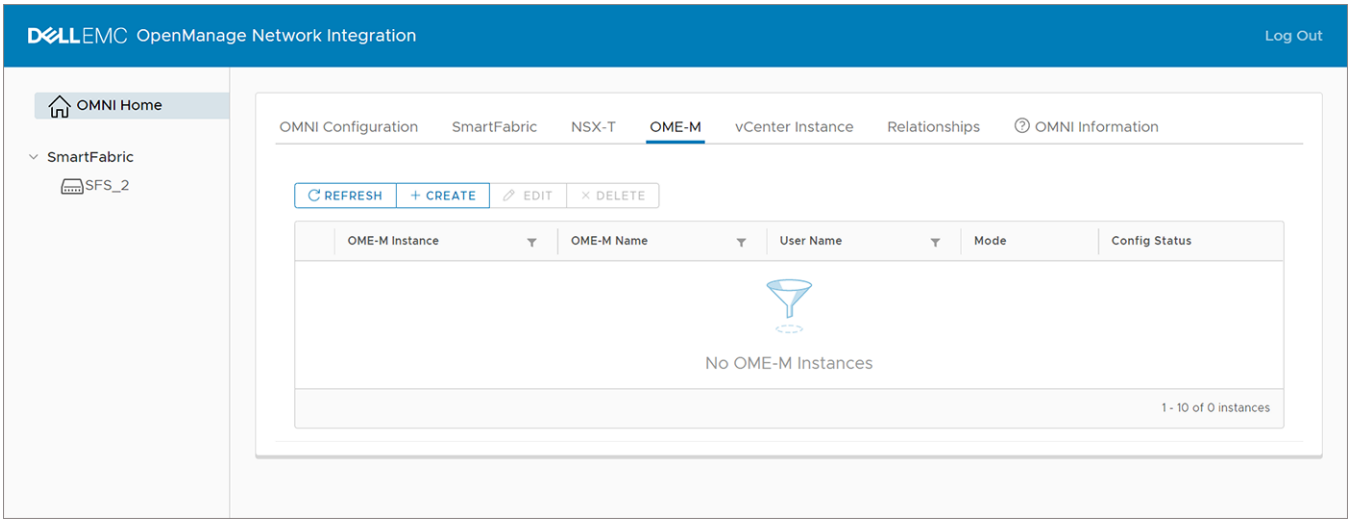

2. Enter the OME-Modular IP address, name, username, and password. Click **Add**.

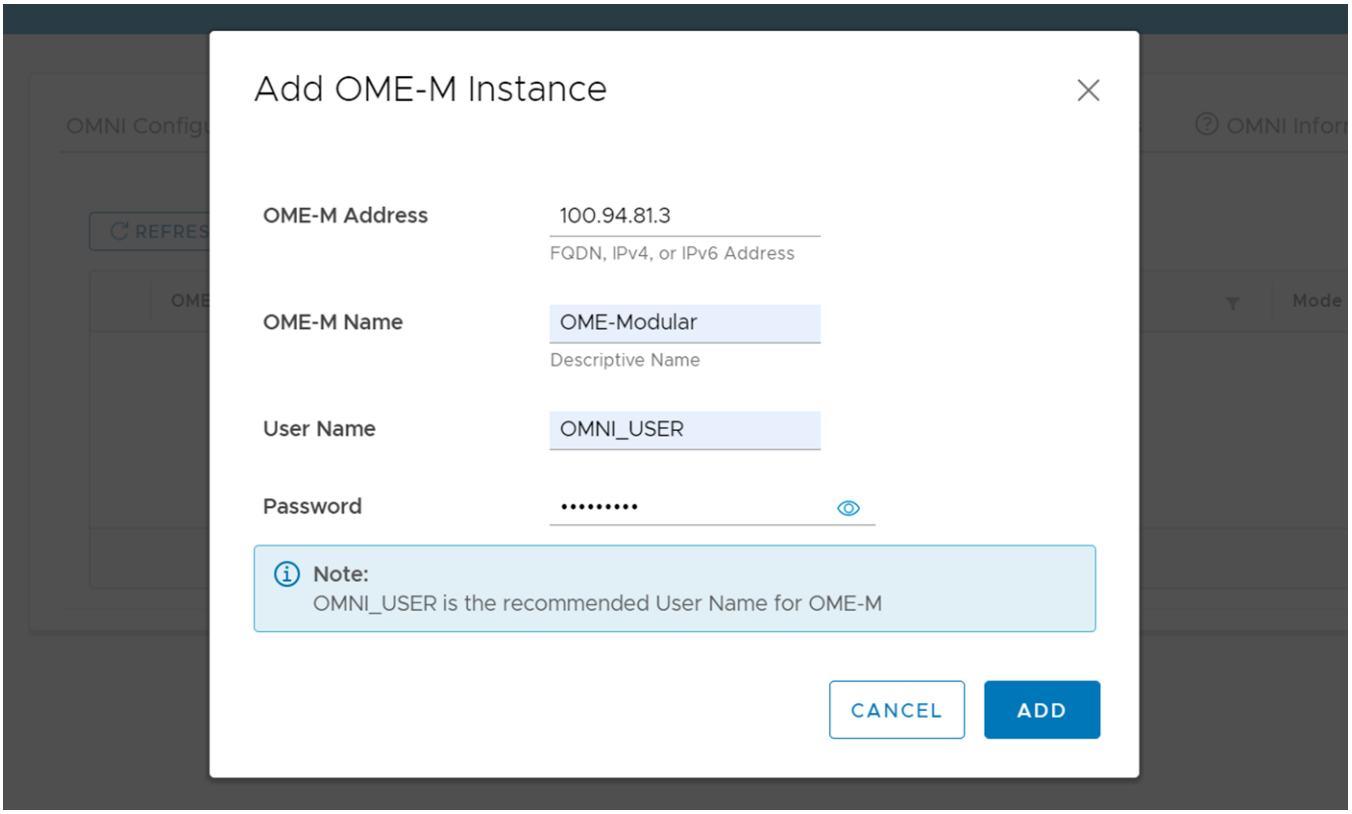

3. The system displays OME-Modular instance creation success message. **OME-M** page displays the list of the service instances available in the OMNI appliance.

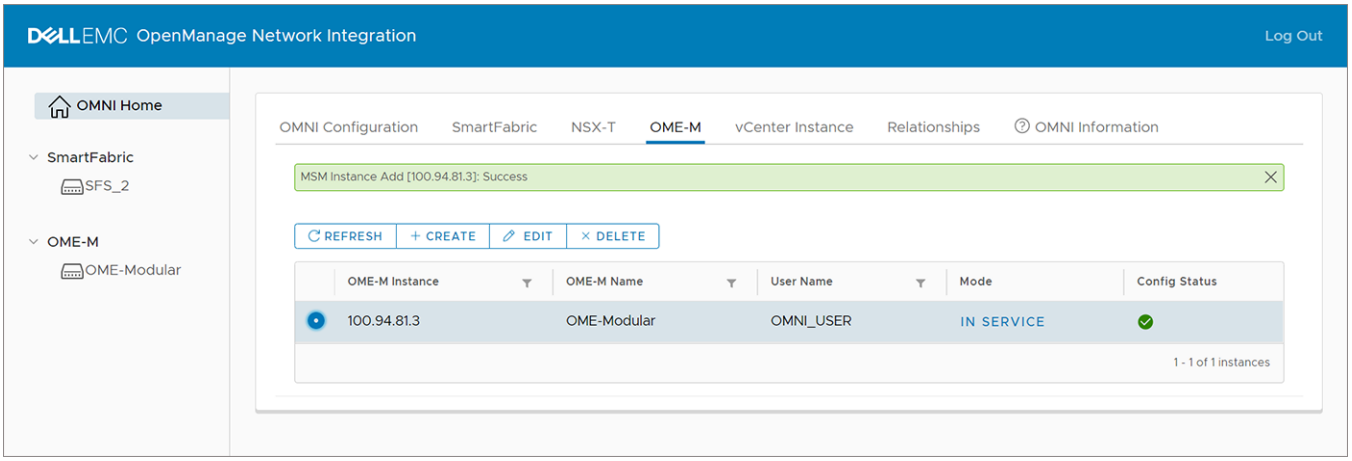

Disable OMNI automation for OME-Modular

Disable automation for the OME-Modular instance by changing the mode from **In Service** to maintenance mode. To do so, click the mode **In Service** and agree to move the mode to **Under Maintenance**. Enabling Maintenance mode stops the automation service for that instance and also disables UI navigation for that instance.

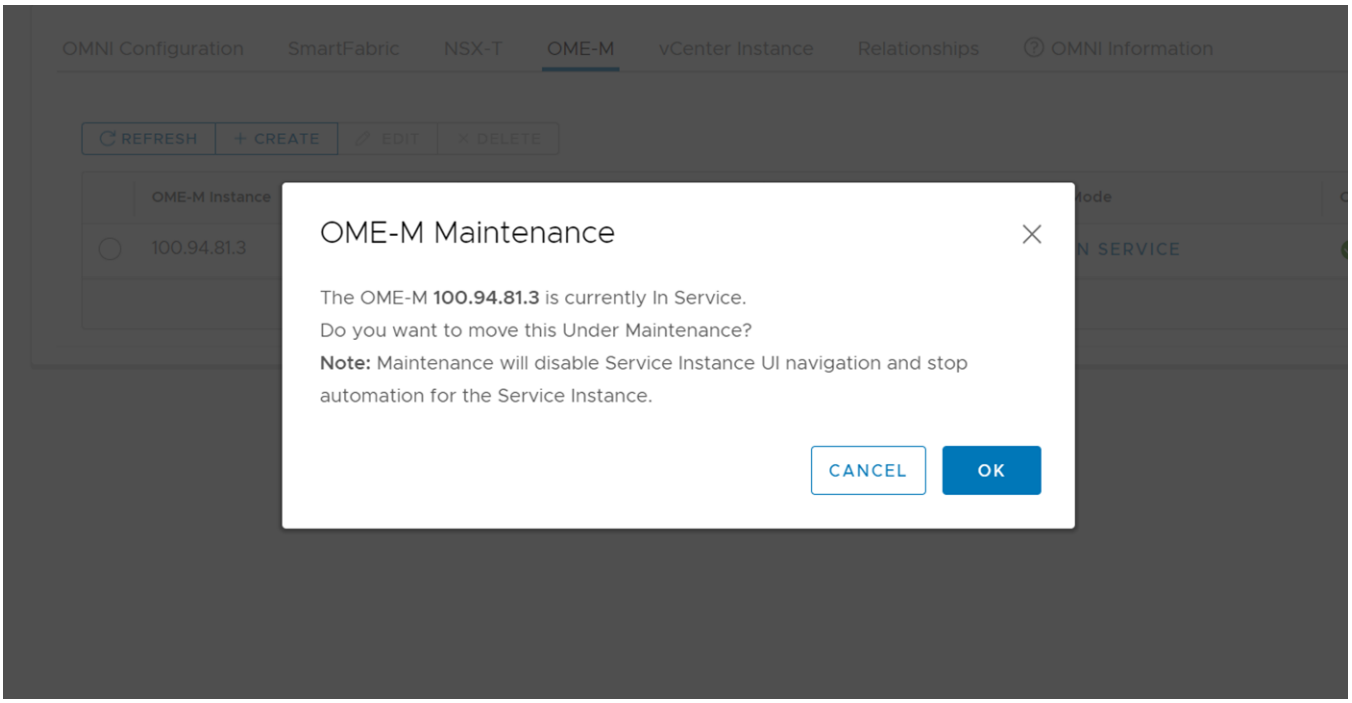

Edit OME-Modular instance

You can edit the name of the OME-Modular instance.

1. From **OMNI Home** > **OME-M**, select the OME-M instance that you want to edit and click **Edit**.

2. Enter the required details and click **Submit**.

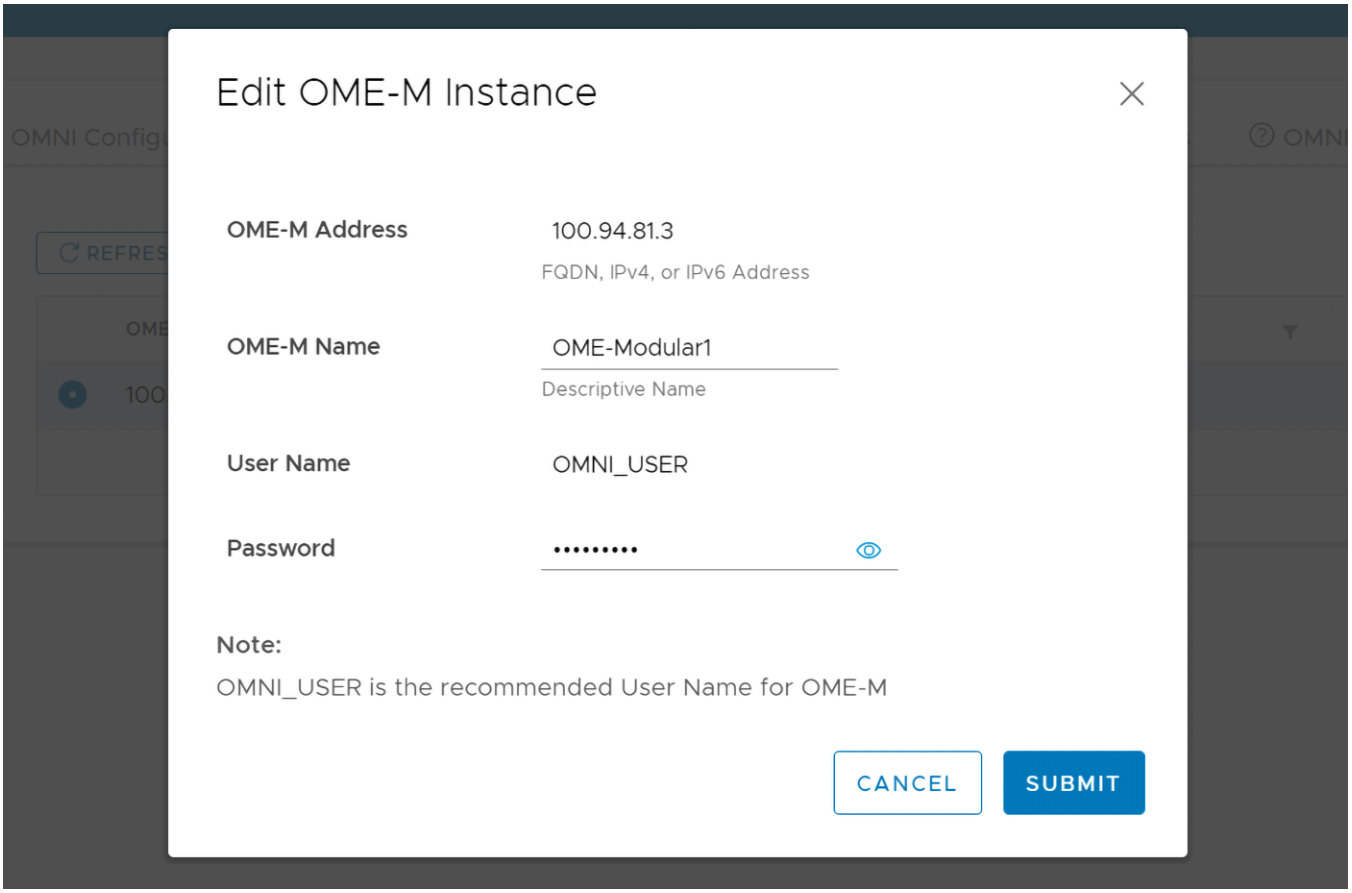

3. The system displays OME-Modular instance edit success message. **OME-M** page displays the list of the service instances available in the OMNI appliance.

Delete OME-Modular instance

You can delete OME-Modular instance from OMNI.

1. Select the OME-Modular instance that you want to delete and click **Delete**.
2. Click **Delete** to confirm.

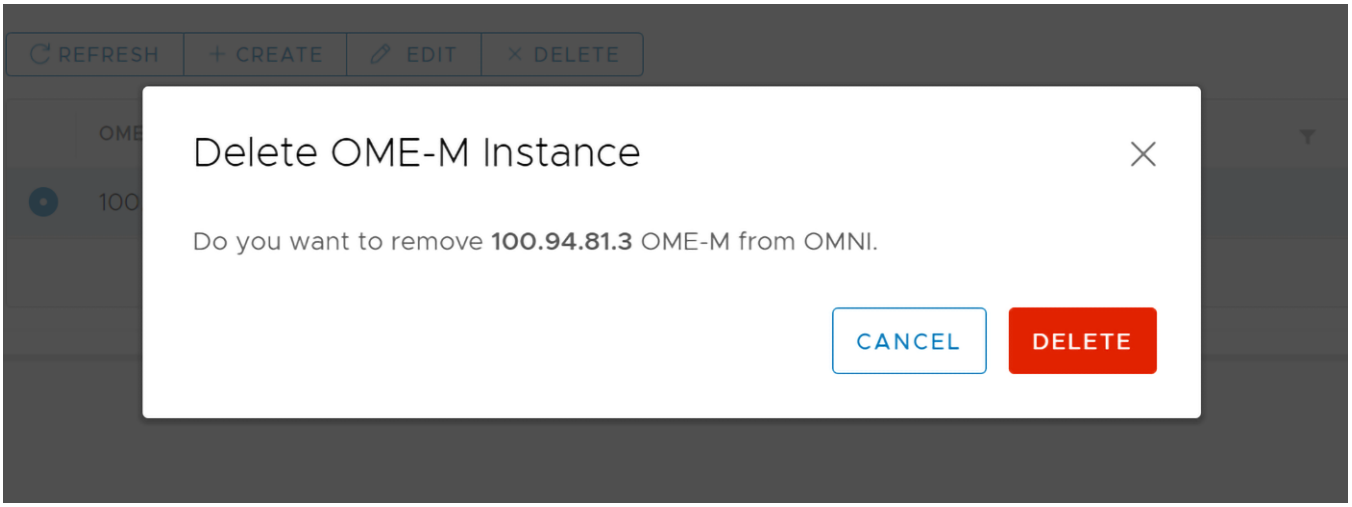

3. The system displays OME-Modular instance delete success message.

Register OMNI with vCenter

Register OMNI with the vCenters that are associated with the MX servers, see Register vCenter with OMNI. When adding the vCenter instance, set the Automation option to **True** to enable the automation services for that vCenter.

View OME-Modular instance

To view the details of the OME-Modular instance, select the OME-Modular instance. OMNI displays a link to OME-Modular console, and you can launch the OME-Modular console by clicking the link.

You can also view the relationship information between the registered vCenter and the OME-Modular. For more information, see View relationship status.

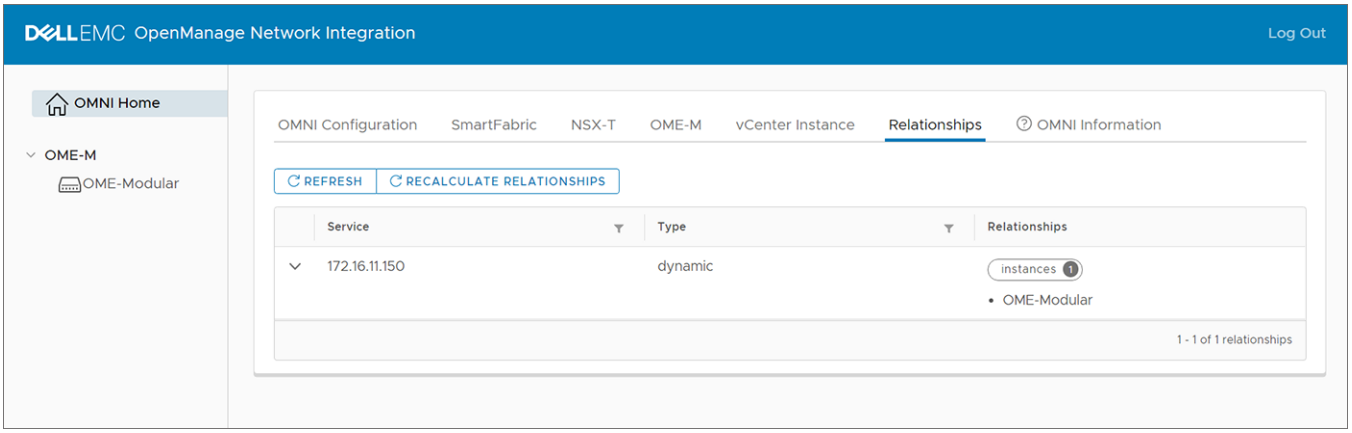

OMNI automation services for OME-Modular

After you add an OME-Modular instance and register the respective vCenters, OMNI creates automation services for the added vCenter instances. You can view the vCenter automation services from OMNI Appliance Management UI. For more information, see here.

NOTE: When you update the MX7000 firmware, Dell Technologies recommends stopping the OMNI automation services for the respective OME-M instance manually. To stop the automation service, select the relevant OME-Modular instance and change the state to Maintenance mode. For more information about disabling automation services, see Register vCenter with OMNI.

NOTE: Note: When the OME-M is not reachable from OMNI, the automation services for the OME-Modular instance must be restarted manually after the network connection is reestablished. OMNI synchronizes the vCenter configuration changes with OME-Modular only after you restart the automation services.

OMNI VLAN automation

Following table lists the OMNI automation actions for various configuration scenarios:

Table 11. OMNI automation scenarios

Configuration scenarios	OMNI automation action
For a port group VLAN creation in vCenter	• OMNI checks if the VLAN is already configured in OME-Modular. If it exists, OMNI uses the existing VLAN. If not, OMNI creates a VLAN in OME-Modular.

Table 11. OMNI automation scenarios (continued)

Configuration scenarios	OMNI automation action
	NOTE: In vCenter, trunk VLAN or private VLAN for port groups is not supported. • OMNI associates the newly created vCenter VLAN to the SmartFabric Ethernet uplinks and related servers. **NOTE:** OMNI does not associate VLANs to the SmartFabric uplinks when there is more than one Ethernet uplink in the SmartFabric. Assign VLANs manually to the uplinks using the OME-Modular UI.
For a port group VLAN deletion from vCenter	OMNI removes the VLAN associated with the related servers, but not from the uplinks. OMNI does not remove the VLAN configuration from OME-Modular.
On deletion of OMNI-created VLAN from OME-Modular	During synchronization, OMNI adds the removed VLANs back to the corresponding servers and uplinks in OME-Modular.
On removal of OMNI-created VLAN from an uplink in OME-Modular	During synchronization, OMNI adds the VLANs back to the corresponding uplink in OME-Modular.

View OMNI-created VLANs in OME-Modular

You can use OME-Modular console to view the configuration changes done by OMNI as part of automation. In OME-Modular, the OMNI-created VLAN has the naming convention of NETWORK_<ID> for the name and OMNI NETWORK_<ID> - DO NOT DELETE for description.

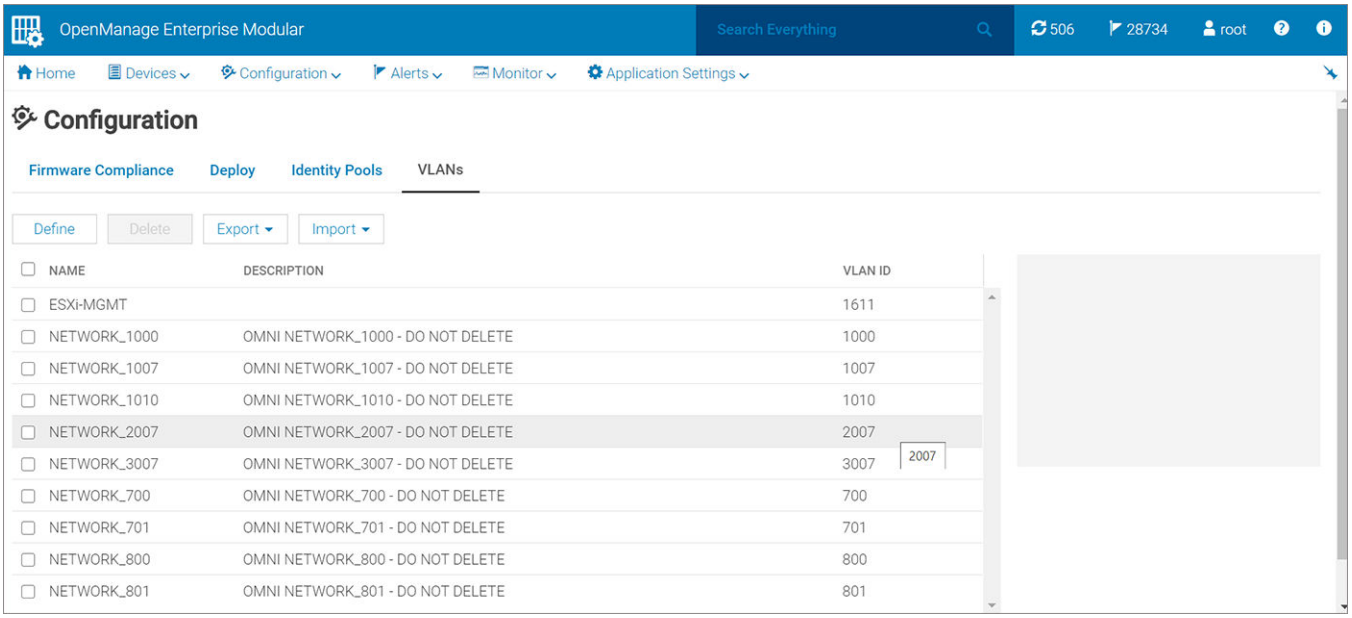

8

OMNI automation support for NSX-T

Starting from 2.0 release, OMNI manages fabric automation for NSX-T. NSX-T is a network virtualization product of VMware that programmatically creates, deletes, and restores software-based virtual networks. For more information about NSX-T, see VMware product documents.

Prerequisites

Ensure that the following prerequisites are met to support OMNI automation services for NSX-T:

- The SmartFabric OS10 version running on the PowerSwtiches should be 10.5.2.2 or a later version that is listed in the SmartFabric OS10 Solutions matrix.
- NSX-T must be in version 3.0.x.
- Servers must be deployed and onboarded in SmartFabric.
- NSX-T Manager cluster must be running and reachable to OMNI.
- The vCenter that is registered with OMNI should be configured as a compute manager in NSX-T.
- OMNI must have connectivity with SmartFabric and vCenter that is registered with OMNI.

Workflow to integrate NSX-T with OMNI

Ensure that the prerequisites are met before starting the workflow to integrate NSX-T with OMNI for automation.

1. Add the NSX-T Manager instance in OMNI.
2. OMNI automation starts for NSX-T.
3. Configure Layer 3 networks that are synchronized from NSX-T.
4. Configure BGP peer routing for edge overlay uplinks.

Add NSX-T instance

To manage the automation for NSX-T using OMNI, add the NSX-T Manager instance in OMNI. In OMNI 2.0 release, you can add one NSX-T Manager instance in a single OMNI VM.

To add a NSX-T instance:

1. From **OMNI Home** > **NSX-T**, click **Create**.

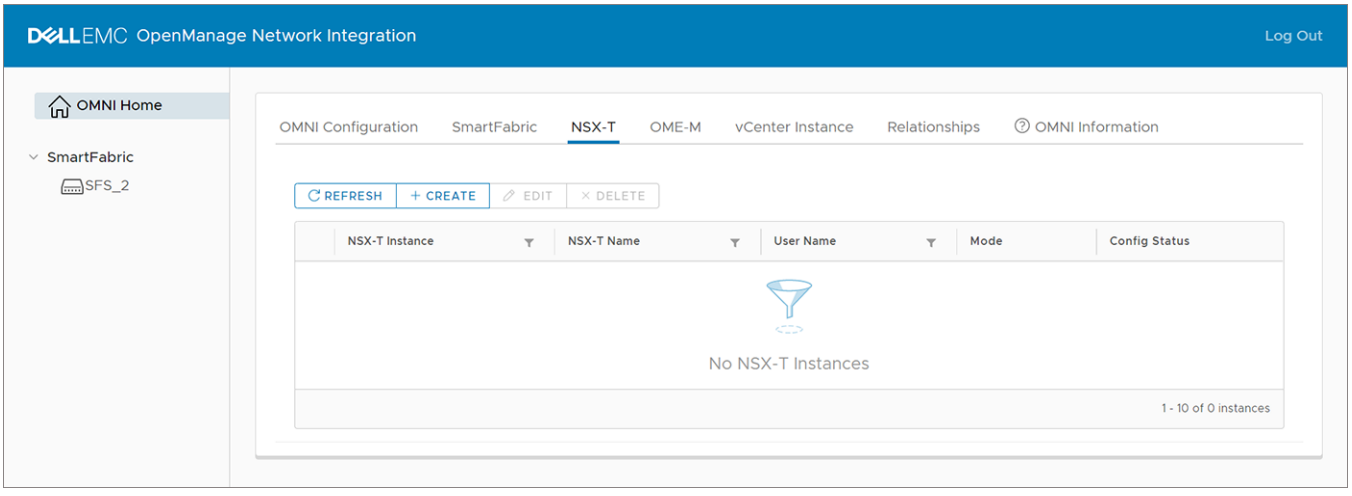

2. Enter the NSX-T Manager cluster virtual IP address or FQDN, name, username, and password. Click **Add**.

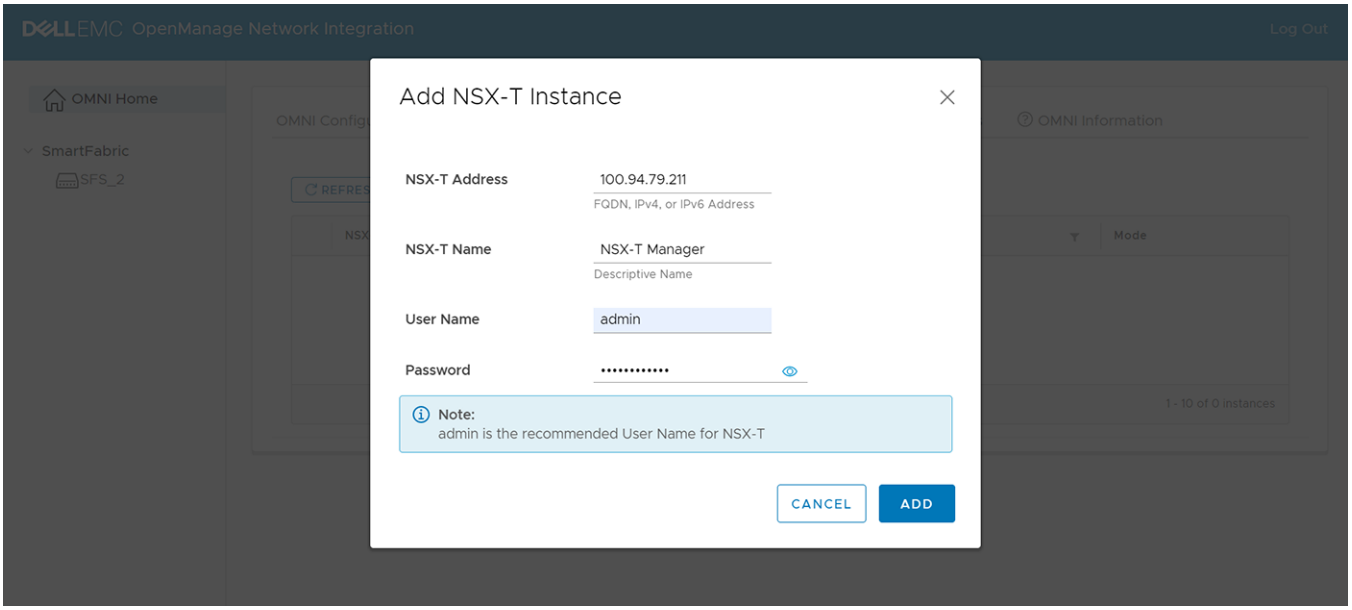

3. The system displays NSX-T instance creation success message. **NSX-T** page displays the list of the instances available in the OMNI appliance.

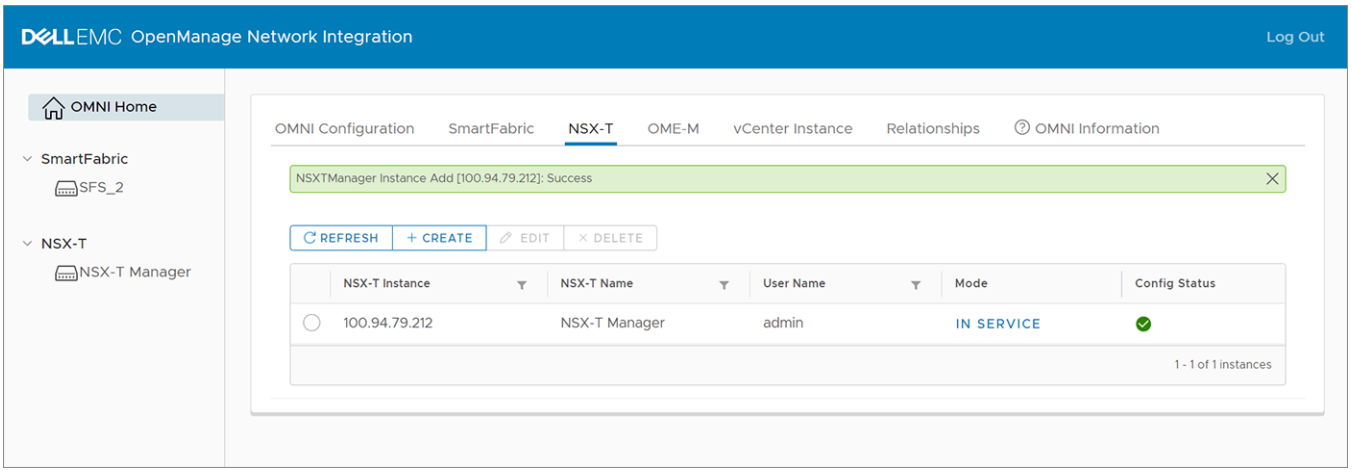

Disable OMNI automation for NSX-T

Disable OMNI automation for the NSX-T instance by changing the mode from **In Service** to maintenance mode. To do so, click the mode **In Service** to change it to **Under Maintenance**. Enabling Maintenance mode disables UI navigation for that instance.

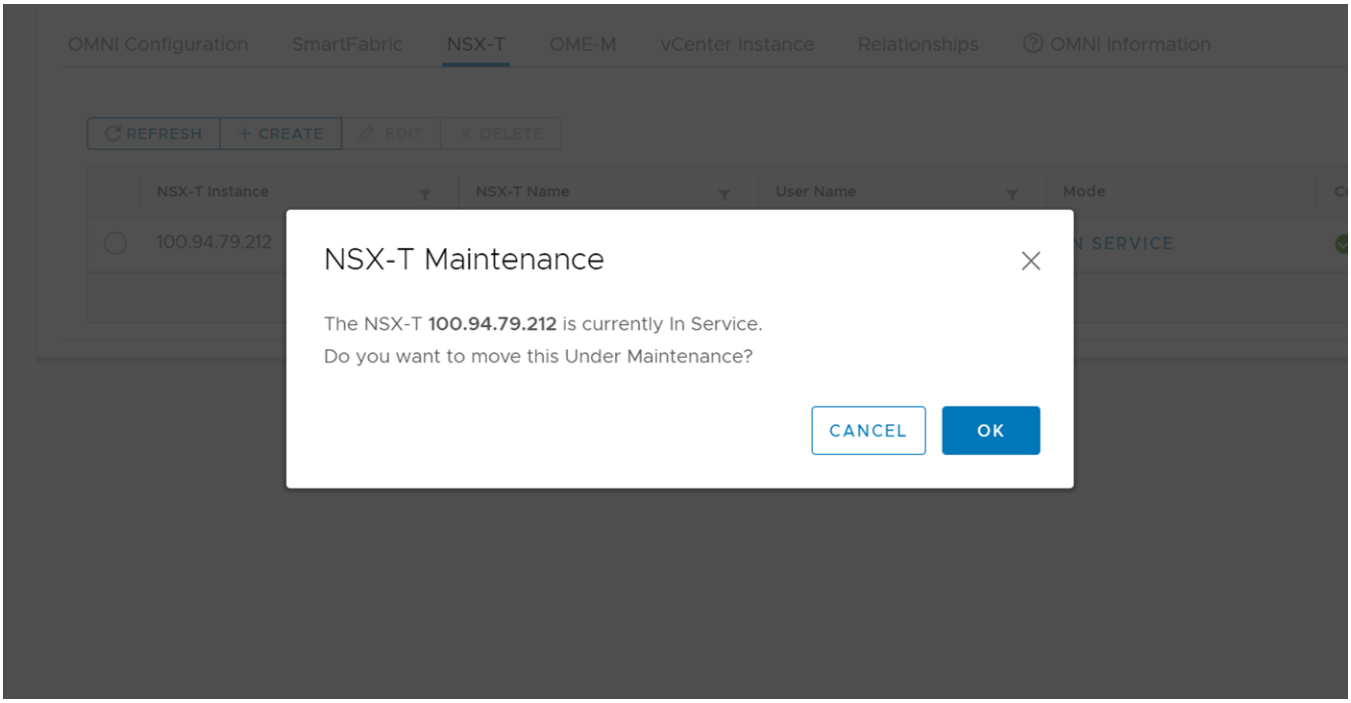

Edit NSX-T instance

You can edit the name of the NSX-T instance.

1. Select the NSX-T instance that you want to edit and click **Edit**.

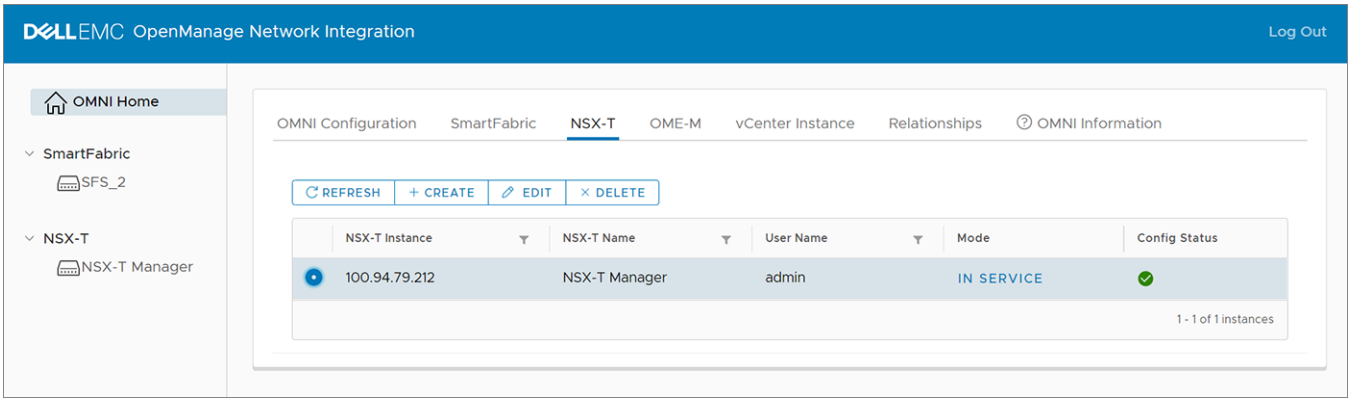

2. Enter the required details and click **Submit**.

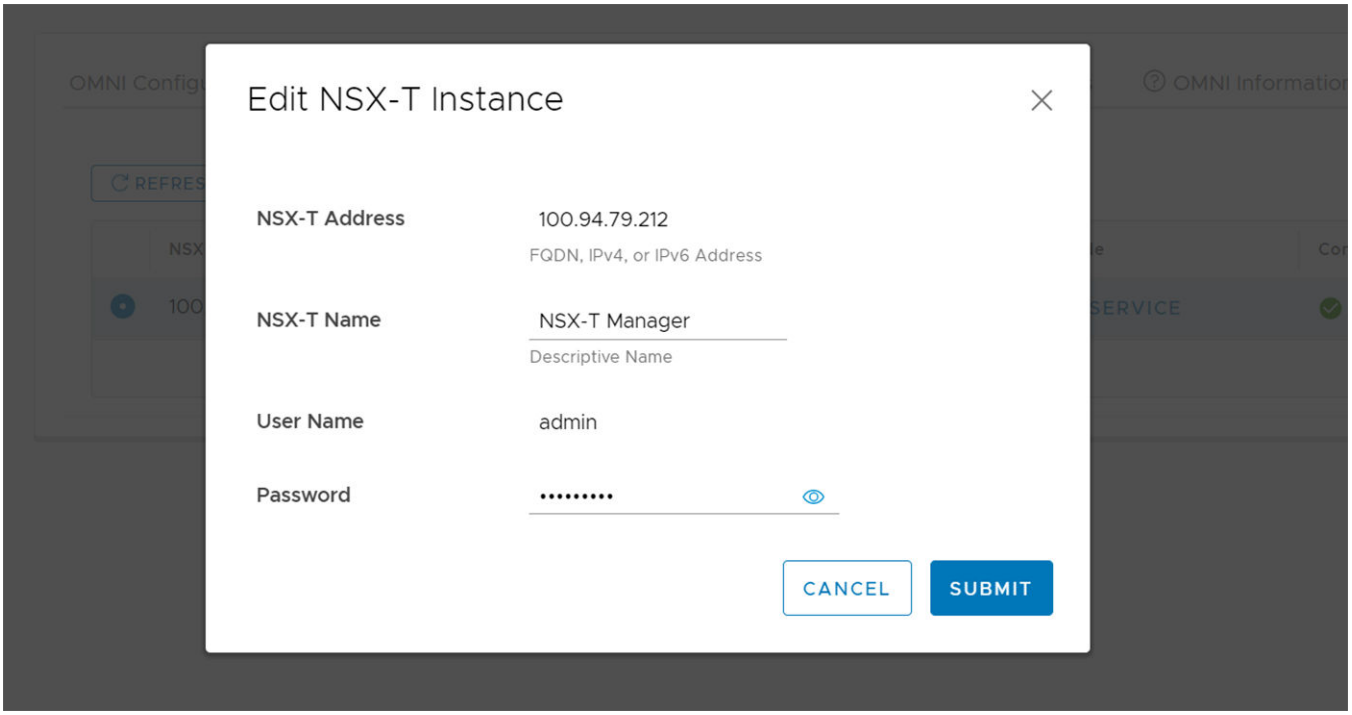

3. The system displays NSX-T instance edit success message. **OME-M** page displays the list of the service instances available in the OMNI appliance.

Delete NSX-T instance

You can delete NSX-T instance from OMNI.

1. Select the NSX-T instance that you want to delete and click **Delete**.
2. Click **Delete** to confirm.

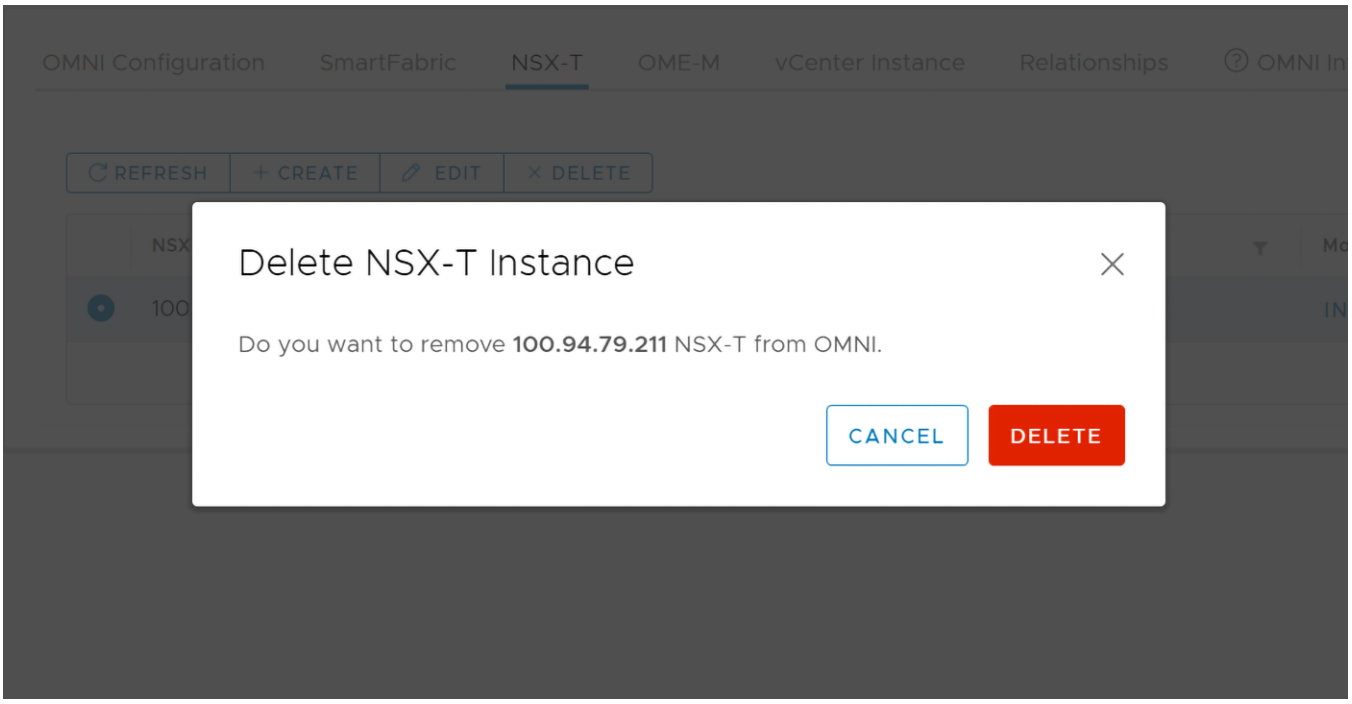

3. The system displays NSX-T instance delete success message.

OMNI automation for NSX-T

After you add the NSX-T Manager as an instance in OMNI, OMNI automation discovers the relationship between the entities such as NSX-T Manager, vCenter, and the SFS instance.

NOTE: It may take few minutes to populate the relationship information and the related networks that are created by OMNI automation.

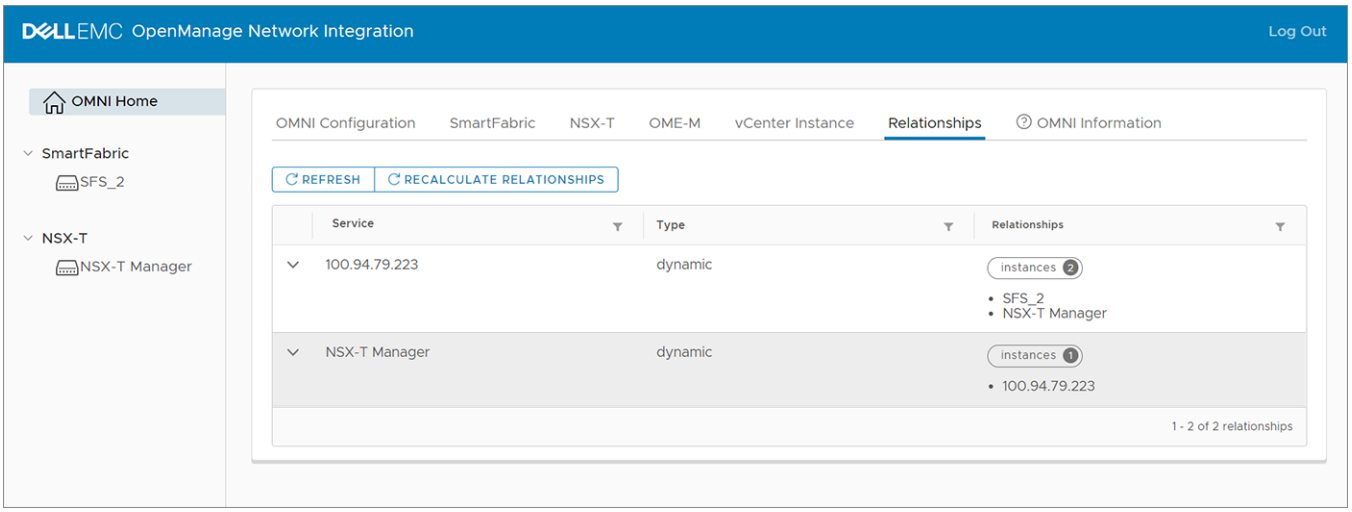

As part of automation, OMNI:

- Creates host and edge overlay networks.
- Creates edge uplink networks.
- Tags the above networks that are created to the corresponding server interface profiles and synchronizes NSX-T networks when you create at least one application segment in NSX-T Manager.

NOTE: OMNI creates all NSX-T networks as multi rack L3 VLAN networks.

OMNI notifies the creation of NSX-T networks using UI alerts.

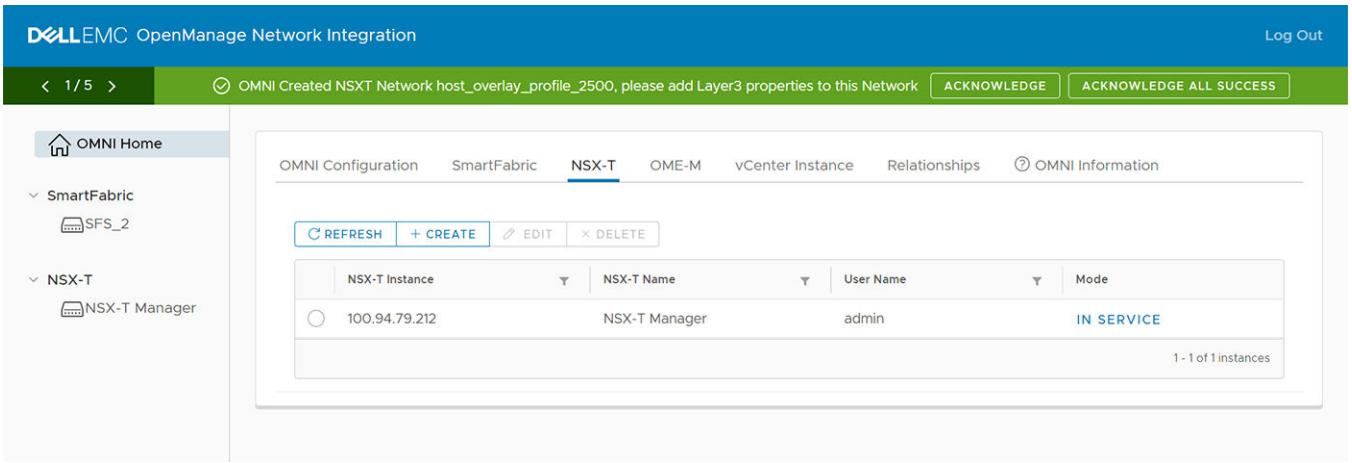

View NSX-T instance

After the NSX-T instance is successfully added, the instance is listed as an entry in the **OMNI Home** left pane.

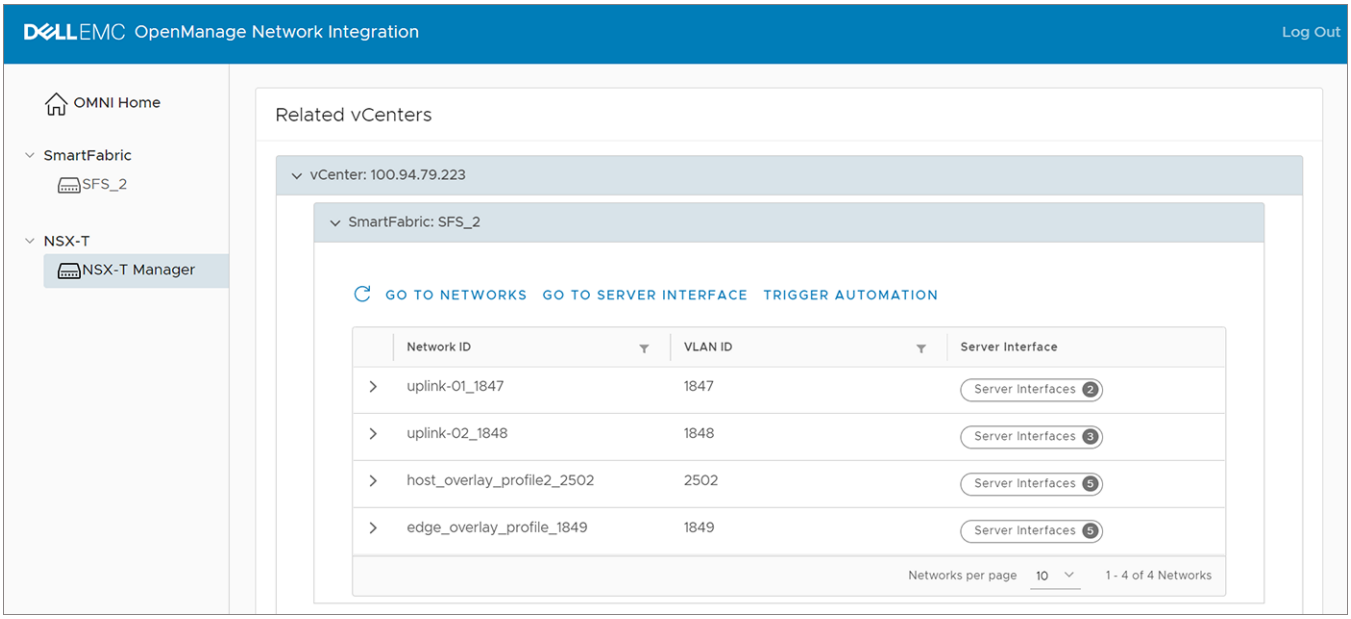

In the left pane, under **OMNI Home**, select the NSX-T instance to view the list of networks that are created by OMNI automation. OMNI displays the vCenter information that is related to the specific NSX-T instance. You can click the vCenter to see the SmartFabric instances that are associated with that instance. When you click a SmartFabric instance, OMNI displays the list of networks that are synchronized from NSX-T and the server interface profiles that are tagged to the NSX-T networks. The NSX-T page displays direct links to **Network** and **Service Interface** configuration tabs of the SmartFabric instance. Click **Trigger automation** to trigger the OMNI automation manually. This action synchronizes the changes in NSX-T to OMNI.

Edit Layer 3 NSX-T networks

After the networks are synchronized from NSX-T, complete the Layer 3 networks configuration in OMNI.

1. From NSX-T instance page, click **Go to Networks**. The click action goes to the **Networks** tab of the SmartFabric instance directly.
2. Click **Multi-Rack L3 VLAN** to see the list of the NSX-T networks that are synchronized by OMNI.

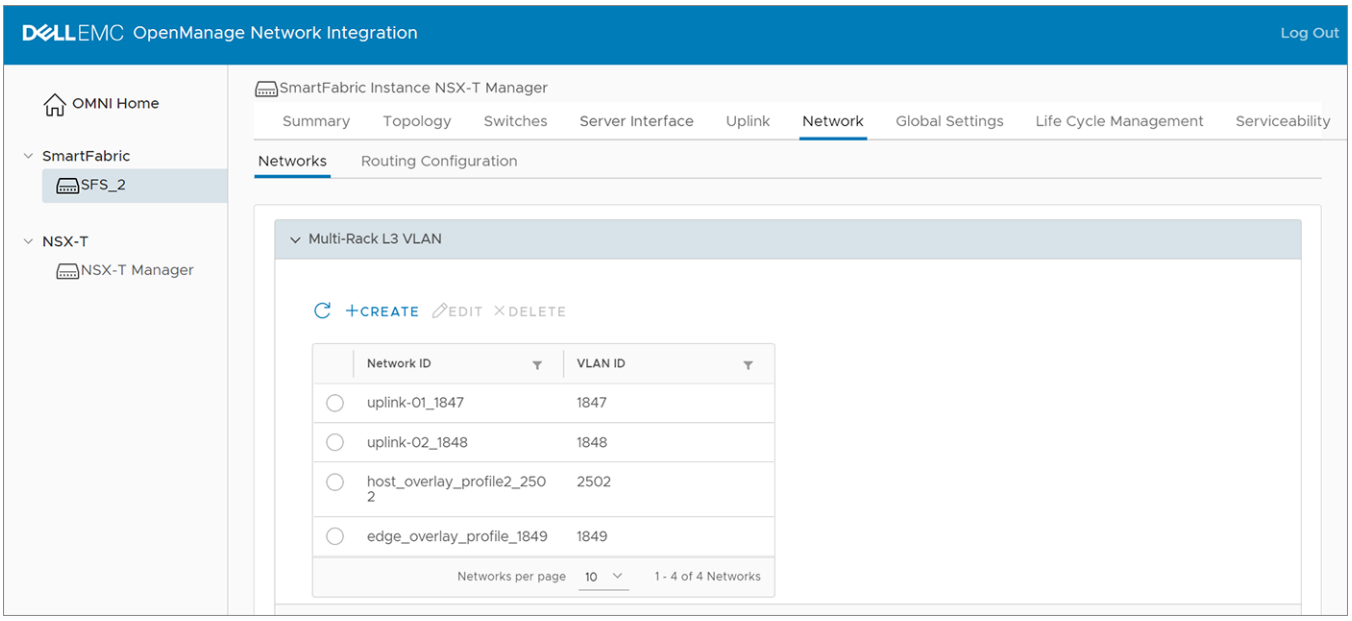

3. Edit all the networks that are synchronized. To edit the Layer 3 settings, see Configure multi-rack L3 VLAN.

Click **Go to Server Interface** to go to the **Server Interface** configuration of the SmartFabric instance. Configure server interface profiles and edit the networks from this page, see server interface profile.

Create edge peer routes

Create edge peer routing IP address details for the uplinks, see Configure eBGP peer route.

9

OMNI support for SmartFabric instances

This chapter explains how to manage SmartFabric components or entities using OMNI. The OMNI VM displays the list of manually created service instances, and the OMNI autodiscovered SmartFabric instances. For more information about the SmartFabric instances, see OMNI Fabric Management Portal.

After you log in to the OMNI Fabric Management Portal, click the SmartFabric instance added to the OMNI Home left page to access and manage the SFS entities that are configured in a SmartFabric.

NOTE: The features that are listed in this chapter are not supported on OME-Modular instance. For more information, see OMNI feature support matrix .

For each SmartFabric instance, you can:

- View the overview of the fabric.
- View fabric topology design.
- Manage switches in a SmartFabric instance.
- Manage server interface configuration.
- Manage uplinks.
- Manage network configuration.
- Configure SmartFabric switch services settings.
- View latest fabric event and compliance errors

OMNI feature support matrix

This table lists the OMNI feature support matrix for SFS-VxRail and PowerEdge MX SmartFabric Services solutions.

Table 12. OMNI feature support matrix for solutions

OMNI feature	SFS-VxRail	PowerEdge MX SmartFabric Services
Service instance and vCenter addition	Yes	Yes
View vCenter and the service instances relationship information	Yes	Yes
vCenter automation	Yes	Yes
Fabric summary	Yes	Managed in OME-Modular
Topology view	Yes	
Switch Inventory	Yes	
Uplink view and creation	Yes	
Network or VLAN view and creation	Yes	
Global configuration	Yes	
SmartFabric OS upgrade	Yes	
Switch replacement in a fabric	Yes	
Fabric backup and restore	Yes	
Fabric event and compliance	Yes	
Server interface profiles view and creation	Yes	NA - use OME-Modular server templates and profiles

View SmartFabric instance overview

Starting from 2.0 release, OMNI displays a consolidated view of key metrics such as device status and health, latest fabric events, and fabric compliance errors for the SmartFabric instances.

From **OMNI Home**, select the SmartFabric instance > **Smmary** > **Overview** to view the dashboard.

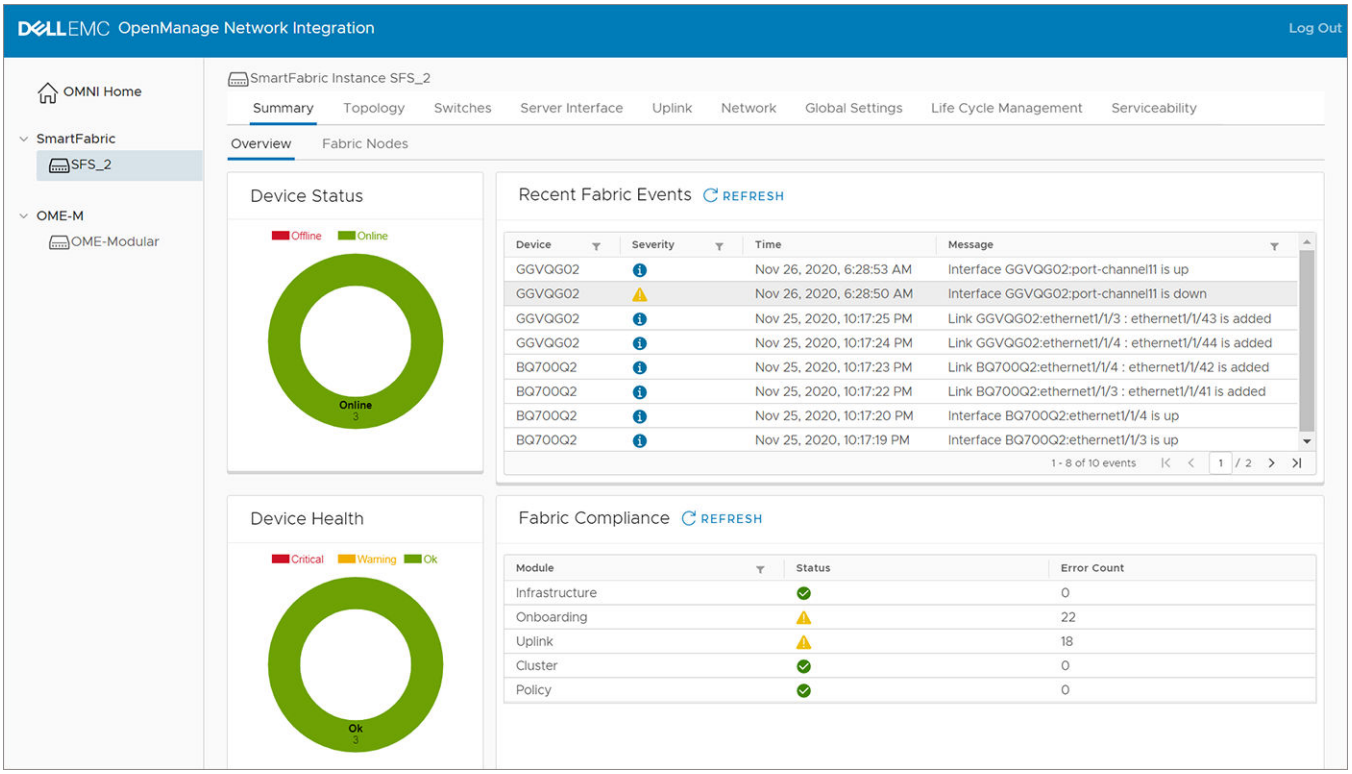

The **Overview** dashboard displays information regarding the following metrics:

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

- Green - Indicates that the devices are online.
- Red - Indicates that the devices are offline.

Recent Fabric Events—Displays the recent fabric events that are generated by SFS. The events are displayed with the following information:

- Device—Service tag of the switch.
- Severity—Severity of the event
  - Critical—Event that is critical that has significant impact.
  - Warning—Event that you should be aware of.
  - Information—Event that does not impact and for informational purpose.
- Time—Time at which the event has occurred.
- Message—Short message about the event occurred.

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

Fabric Compliance—Displays the misconfiguration and compliance violations that are identified in the instance by SFS.

- 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 events in the SmartFabric instance from **Serviceability** page.

View node details

To view the details of the nodes or switches in the fabric, select the SmartFabric instance > **Summary** > **Fabric Nodes**.

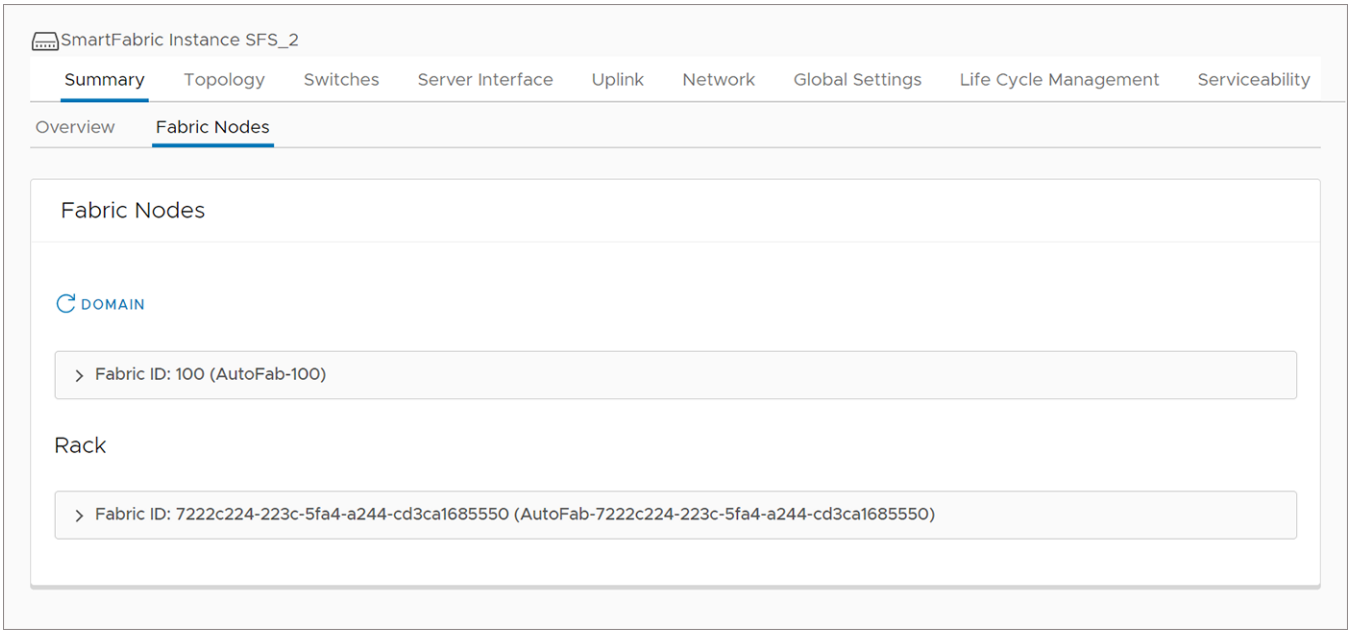

From **Fabric Nodes**, view the list of spine and leaf nodes that are deployed in the fabric. Each switch includes status (online or offline), name, model, version, role, and IP address.

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

Fabric ID—Displays the status of spine switches connected in the fabric.

Rack—Displays the status of the leaf switches in each rack.

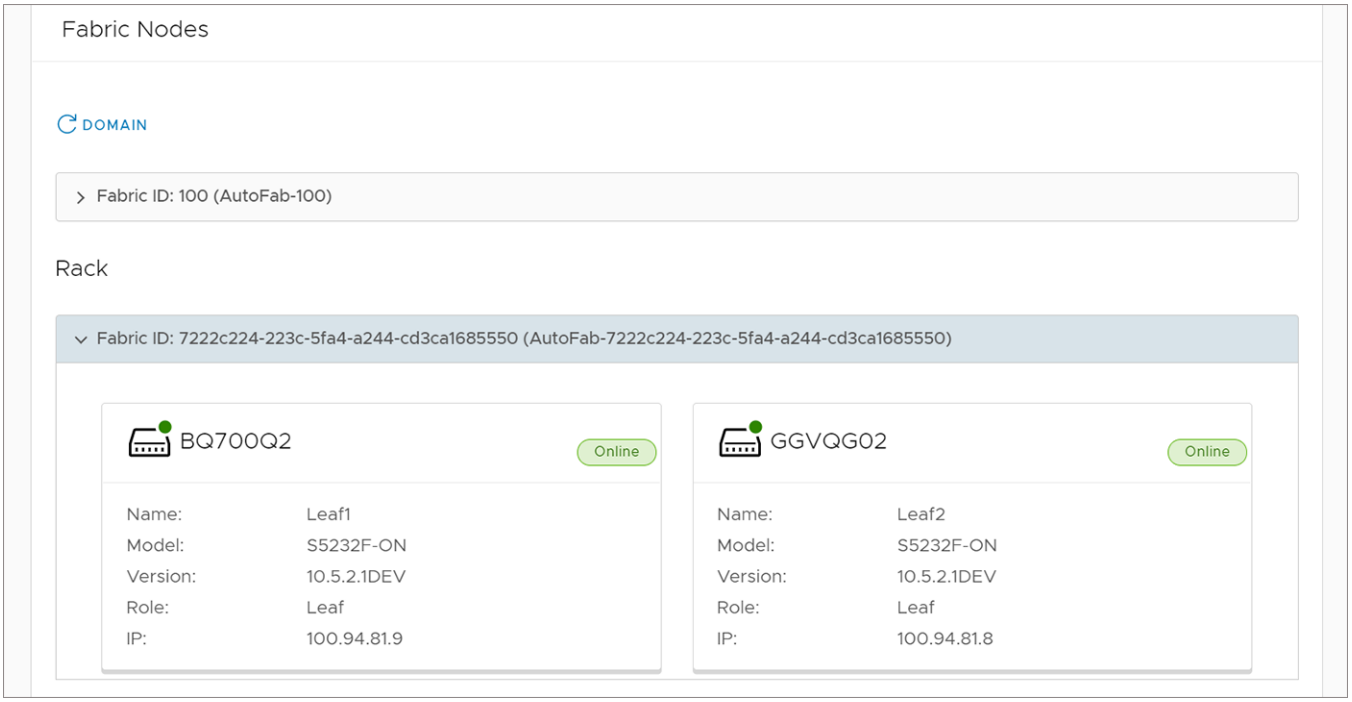

View fabric topology

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.

From the SmartFabric instance, select **Topology** to view the graphical representation of the fabric.

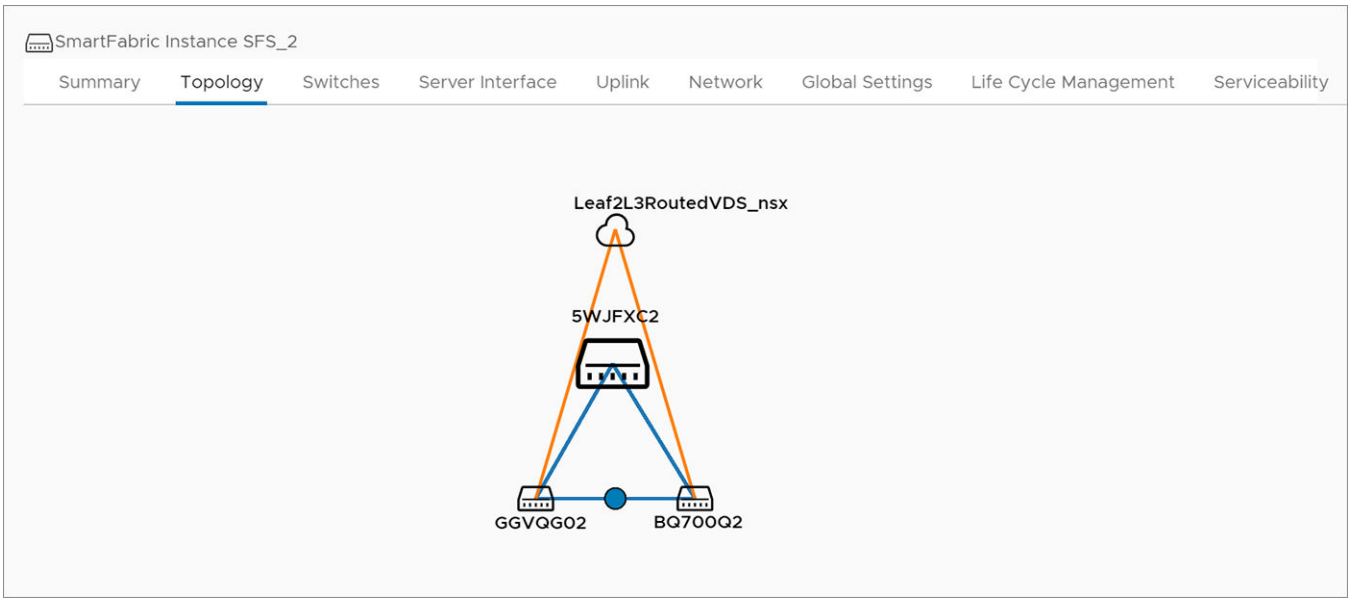

The topology view displays the graphical icons of all the nodes and the link connectivity between the nodes. Each graphical node is represented with their service tag. Hover over an icon to view the detailed information about the node, and the link connectivity in the nodes. The detailed information of the node includes switch ID, switch platform, type of switch (leaf or spine), OS10 version running on the switch, and IP address. You can also view the details of source and destination interfaces of the link, when you hover over the links between the nodes.

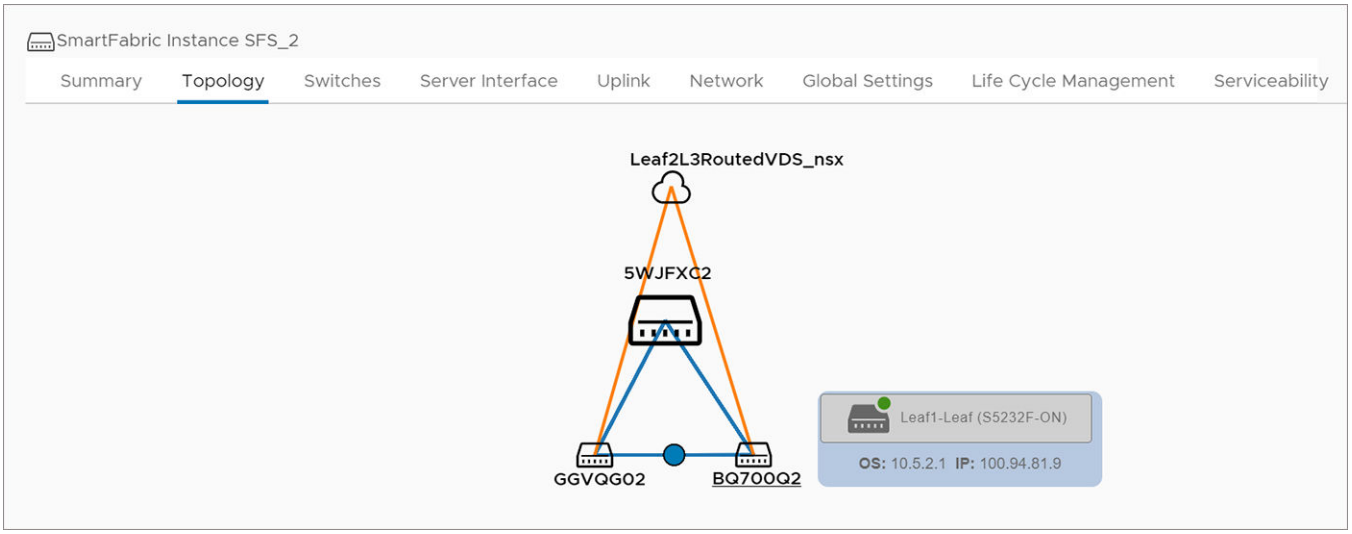

Manage switches in a fabric

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

From **Switches** page:

- View the details of the switches and the ports in a fabric.
- Edit the interface details.
- Set the MTU value for the port.
- Manage the unused ports in the switches.
- Configure breakout ports in leaf switches.
- Configure jump port.

View switch and port details

View the details of the leaf and spine switches, and the list of all ports and unused ports available in each switch. All ports category contains the list of interface and port channel in the switch.

1. From the SmartFabric instance, select **Switches**.

   Fabric Switches—Displays the list of spine and leaf switches available in the selected SmartFabric.

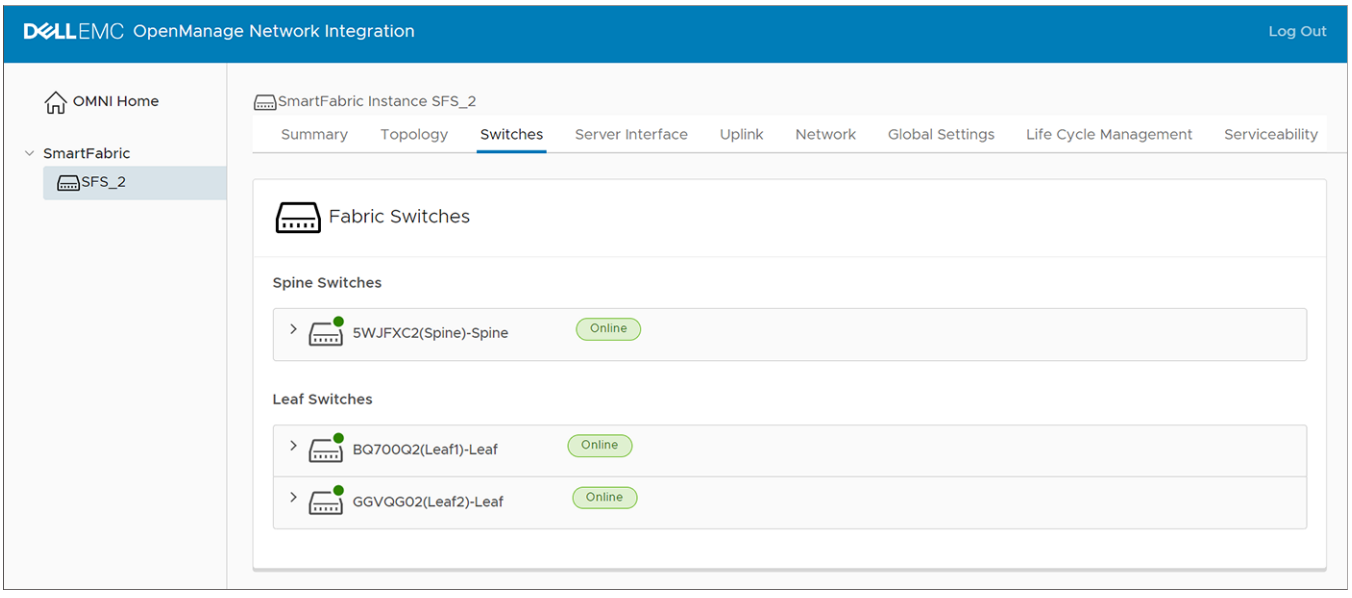

2. Select the arrow of the respective leaf or spine switch to view more information.

Spine Switches—Displays the list of all spine switches with ports information in categories. Click the arrow of the respective switch and category to view more about port information.

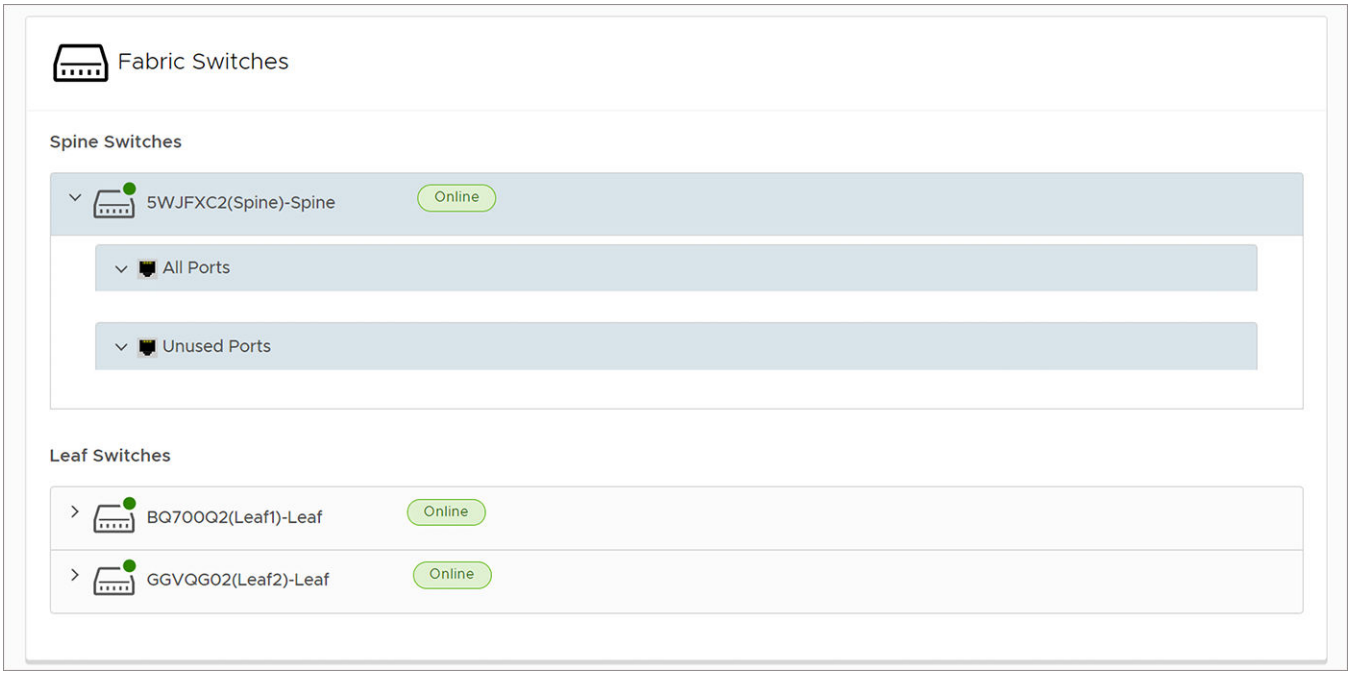

Leaf Switches—Displays the list of all leafs in the fabric with ports, unused ports, breakout ports, and jump port information in categories. Click the arrow of the respective leaf switch category to view more information about the ports.

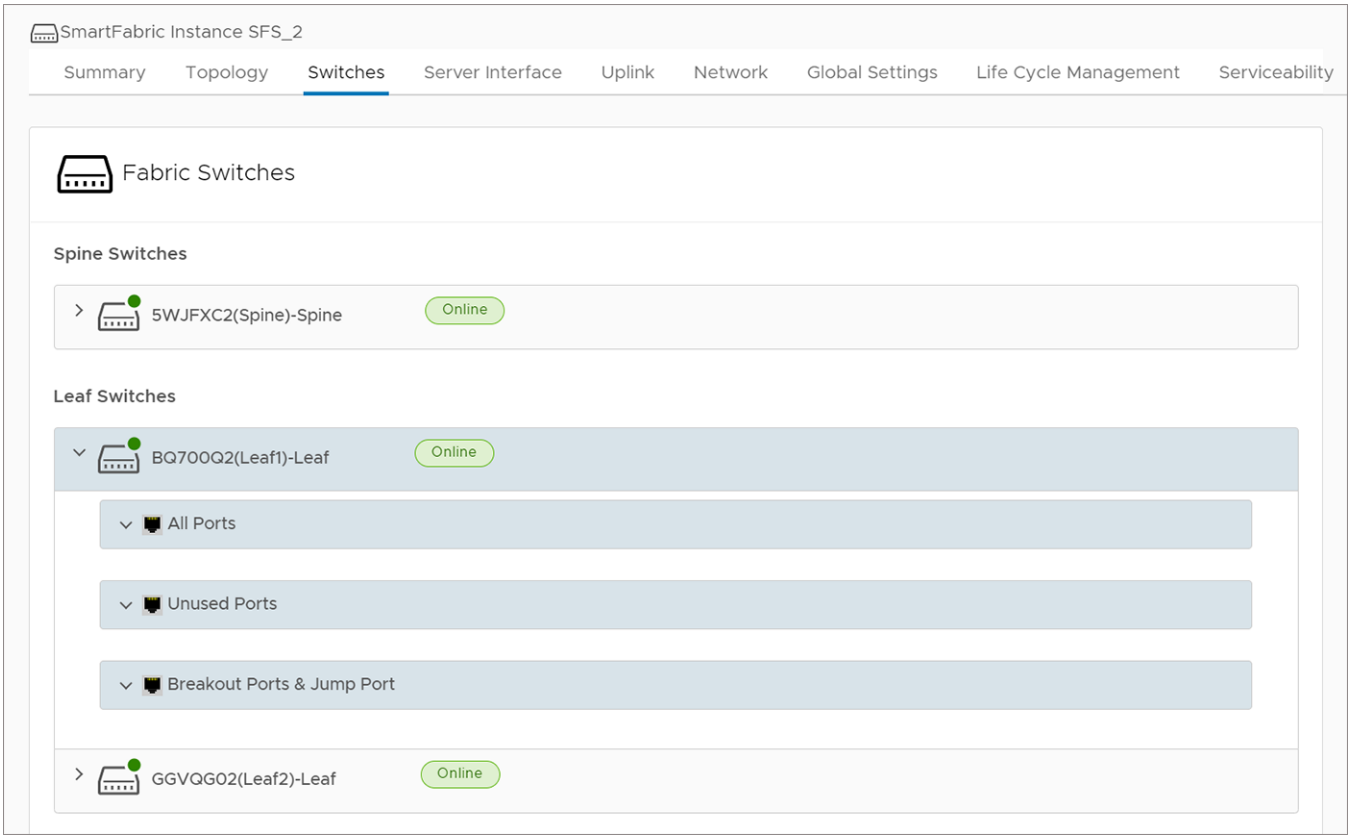

Edit port configuration on a switch

Edit the configuration of port on a leaf or spine switch.

1. From the SmartFabric instance, select **Switches**.

2. Select the spine or leaf switch by clicking the arrow to view more information.

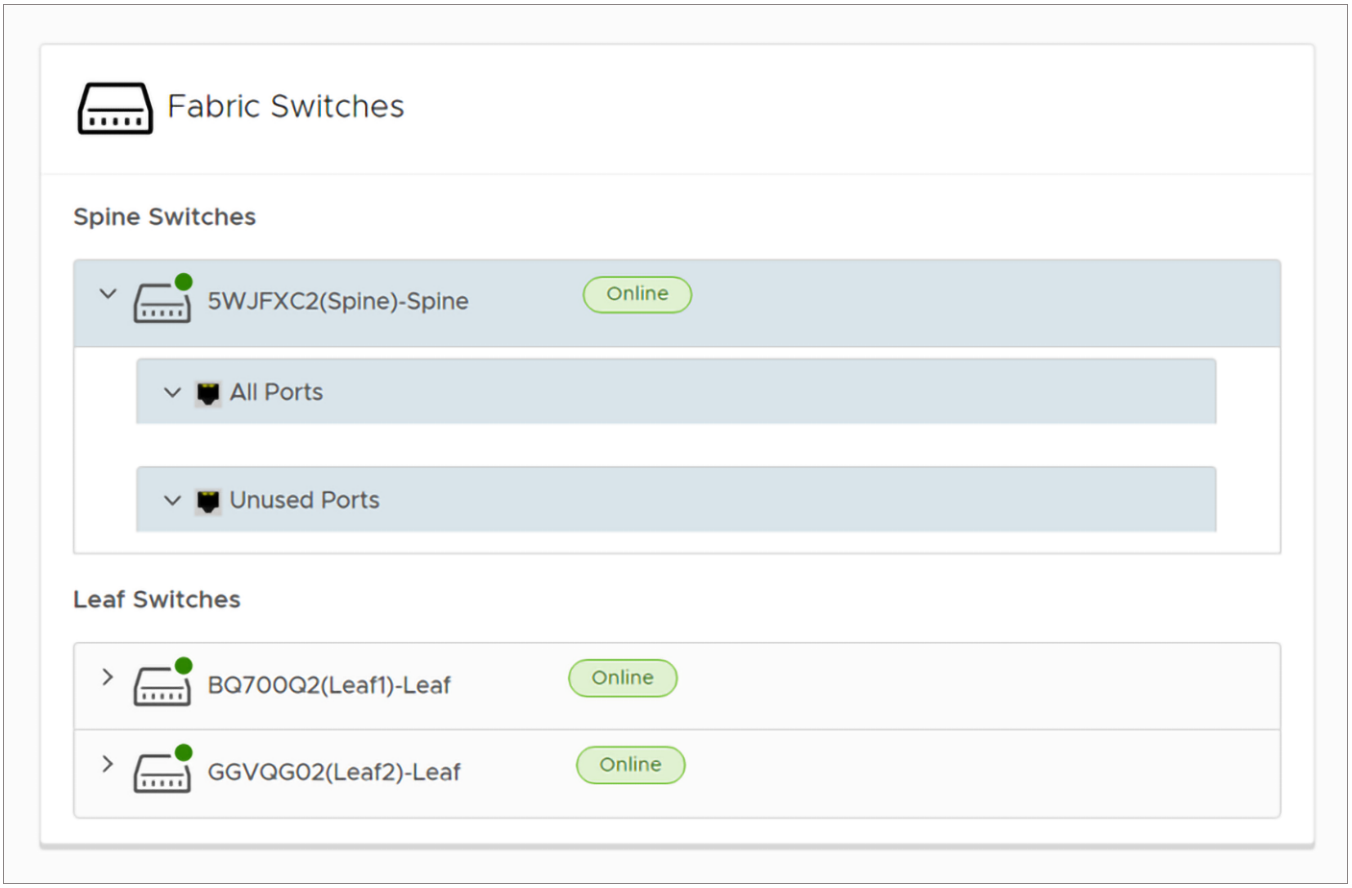

3. Select a port from **All Ports** category, and click **Edit**.

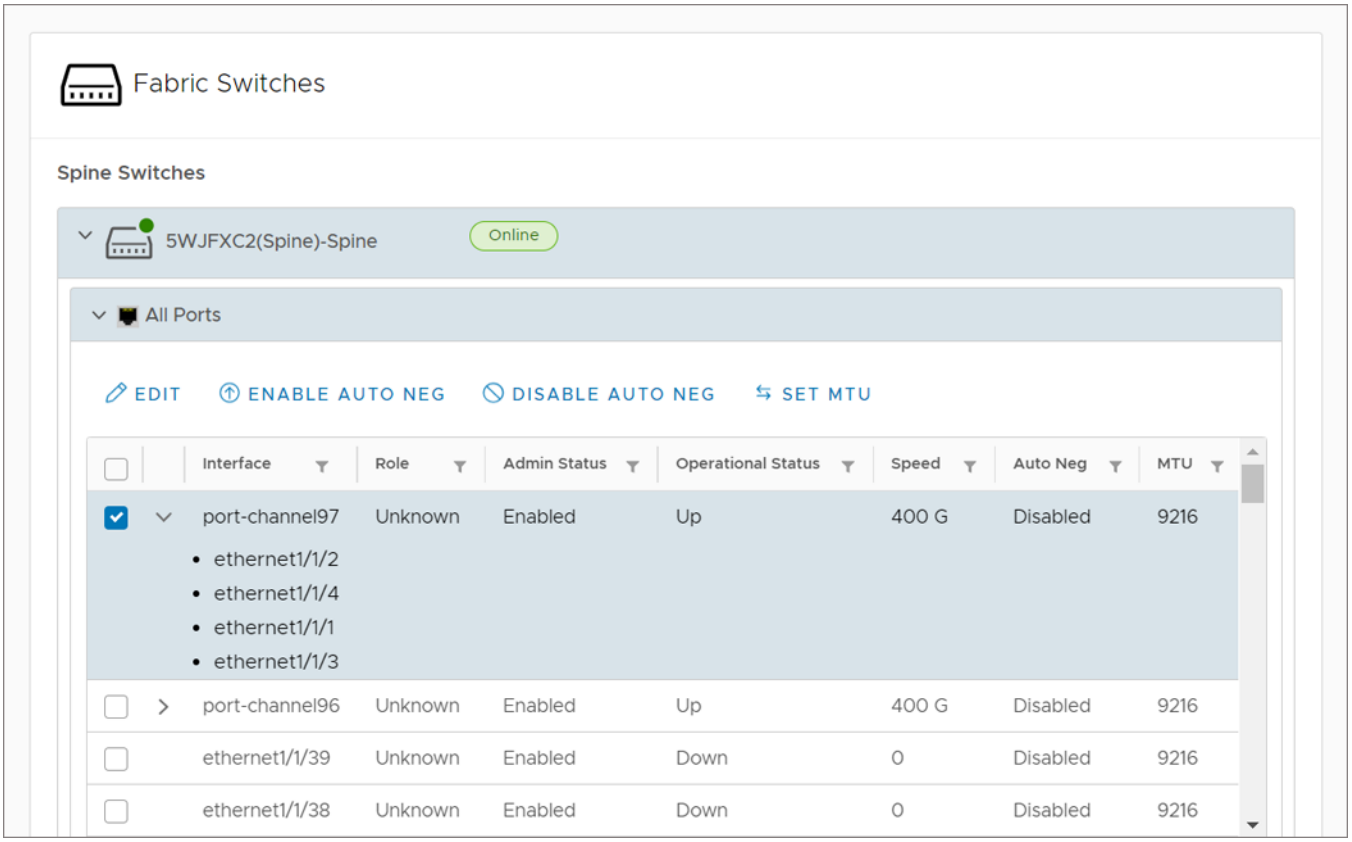

4. Edit the name, description, admin status, auto negotiation, and MTU, and click **Edit**.

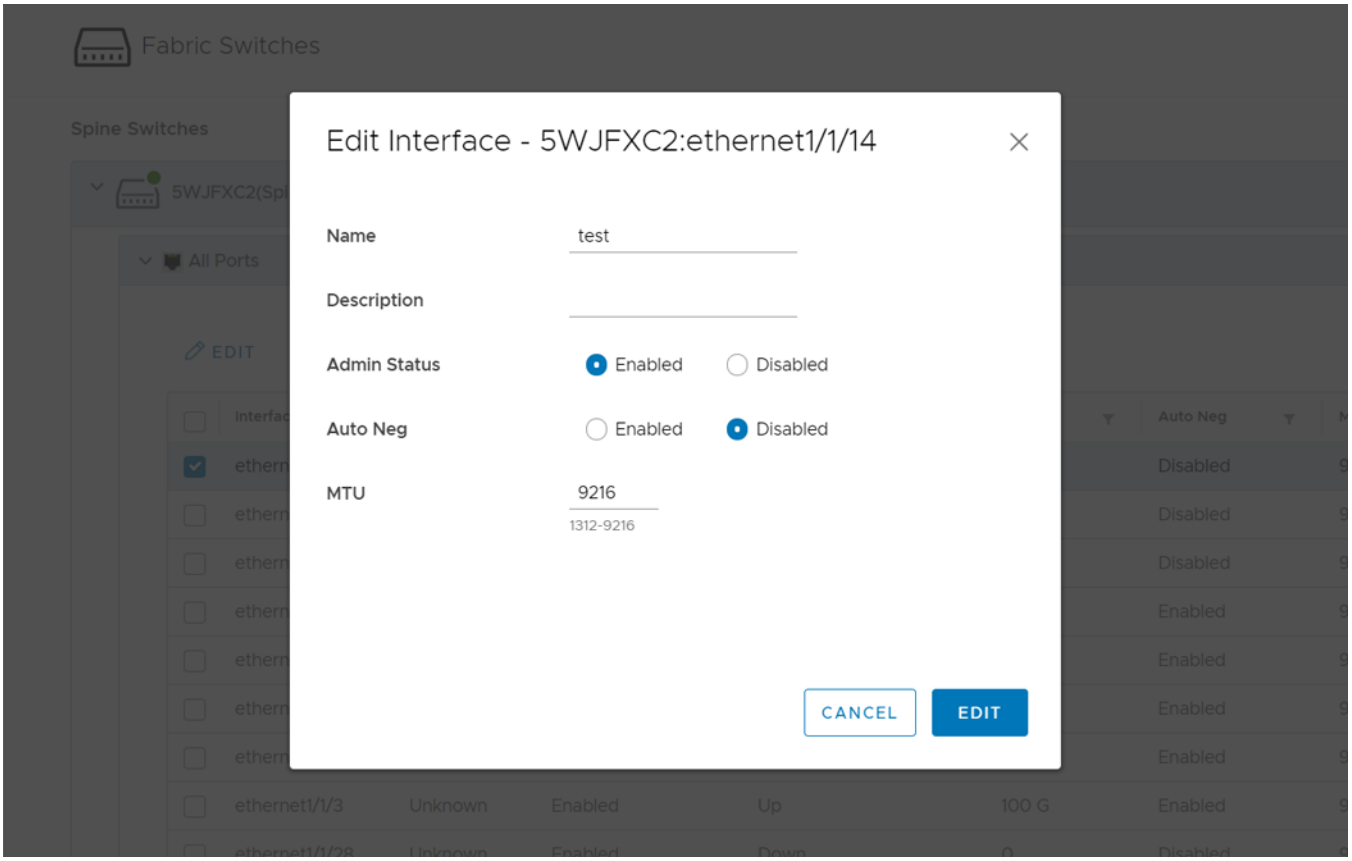

Configure auto negotiation status

You can enable or disable the auto negotiation on a single port or multiple ports.

To enable auto negotiation:

1. From **All Ports**, select a port or multiple ports and click **Enable Auto Neg**.

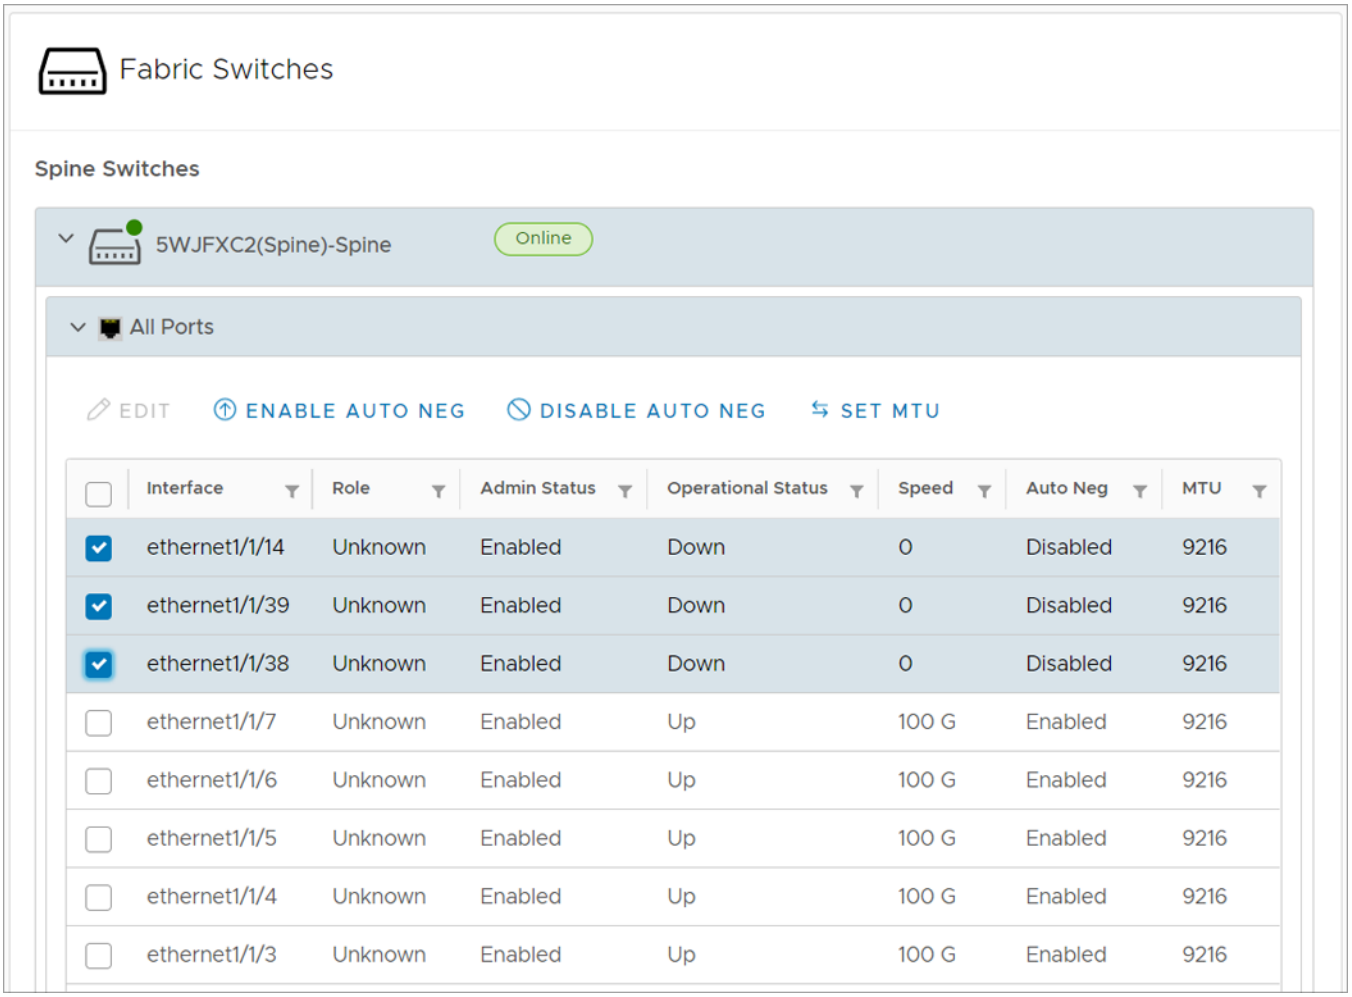

2. The system displays a warning message. Click **Yes** to confirm.

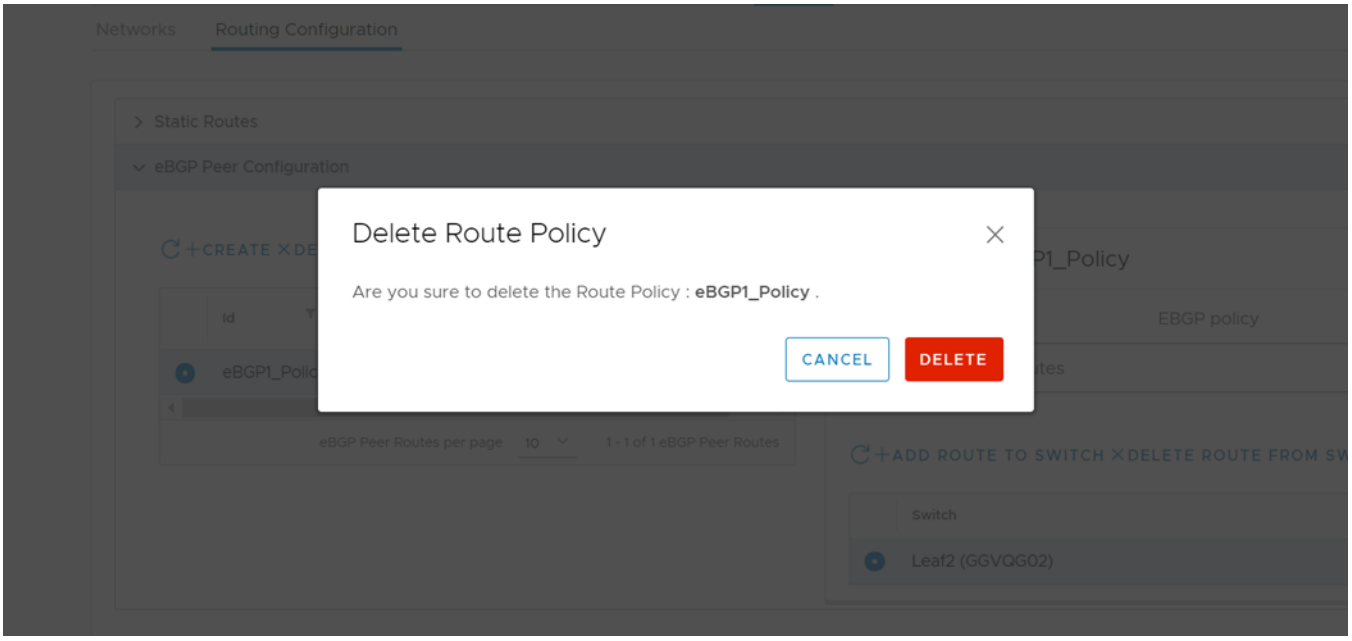

3. The system displays the stage-wise progress of the interface status.

To disable auto negotiation:

1. From **All Ports**, select a port or multiple ports and click **Disable Auto Neg**.
2. 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**.

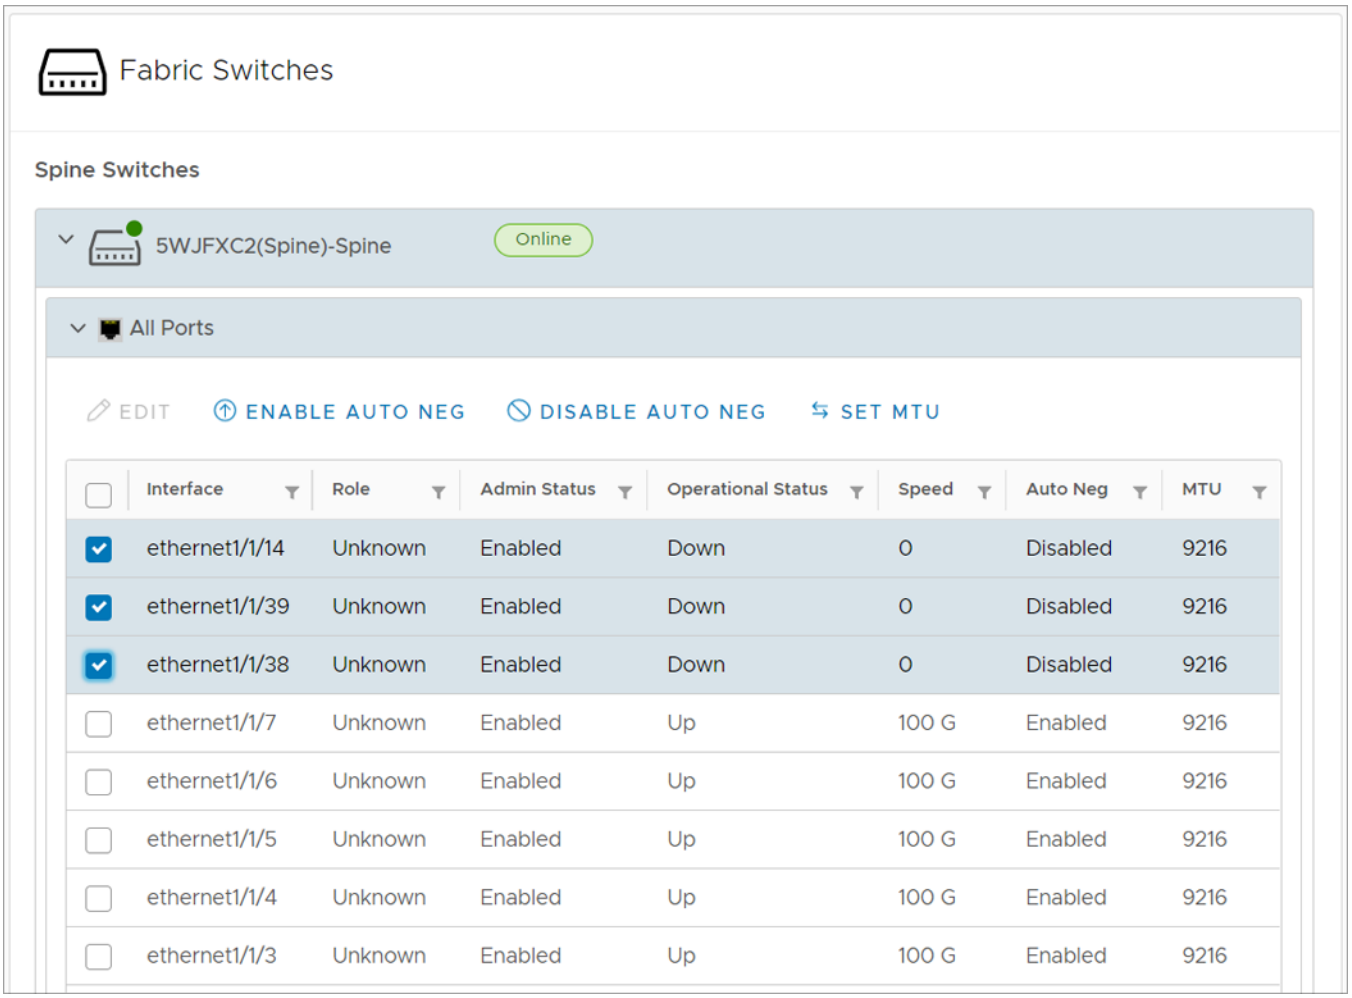

2. Enter the MTU value and click **Set**.

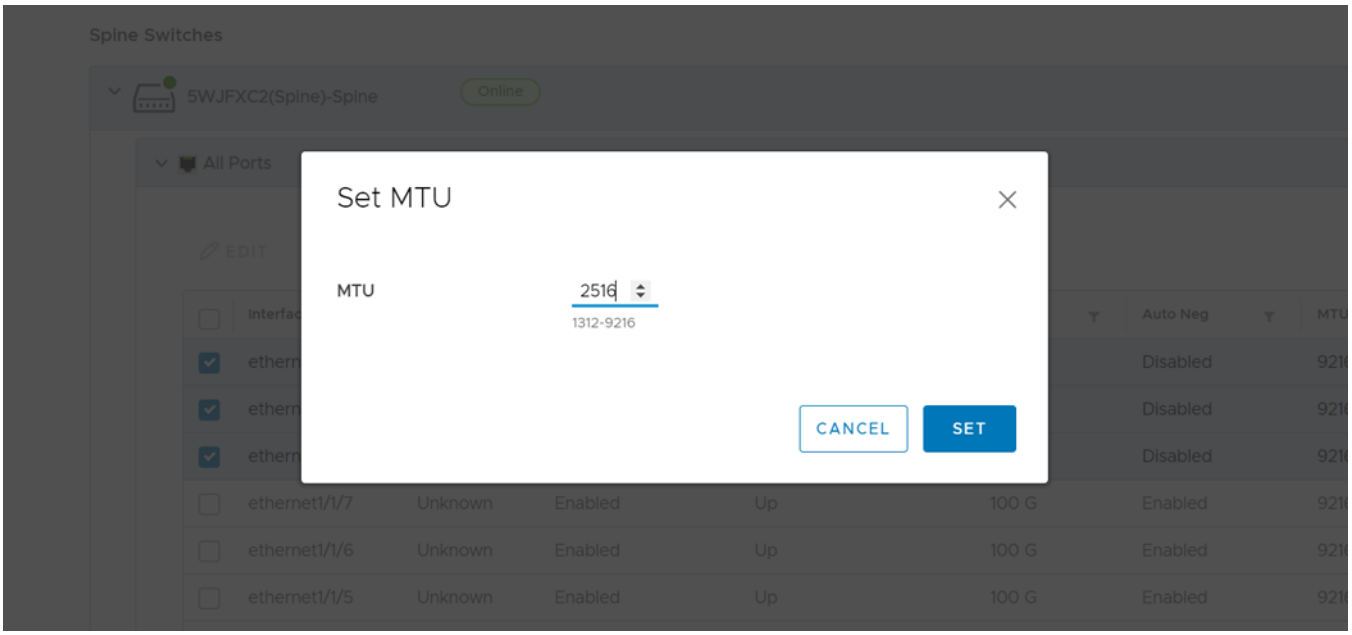

3. The system displays a warning message. Click **Yes** to confirm.

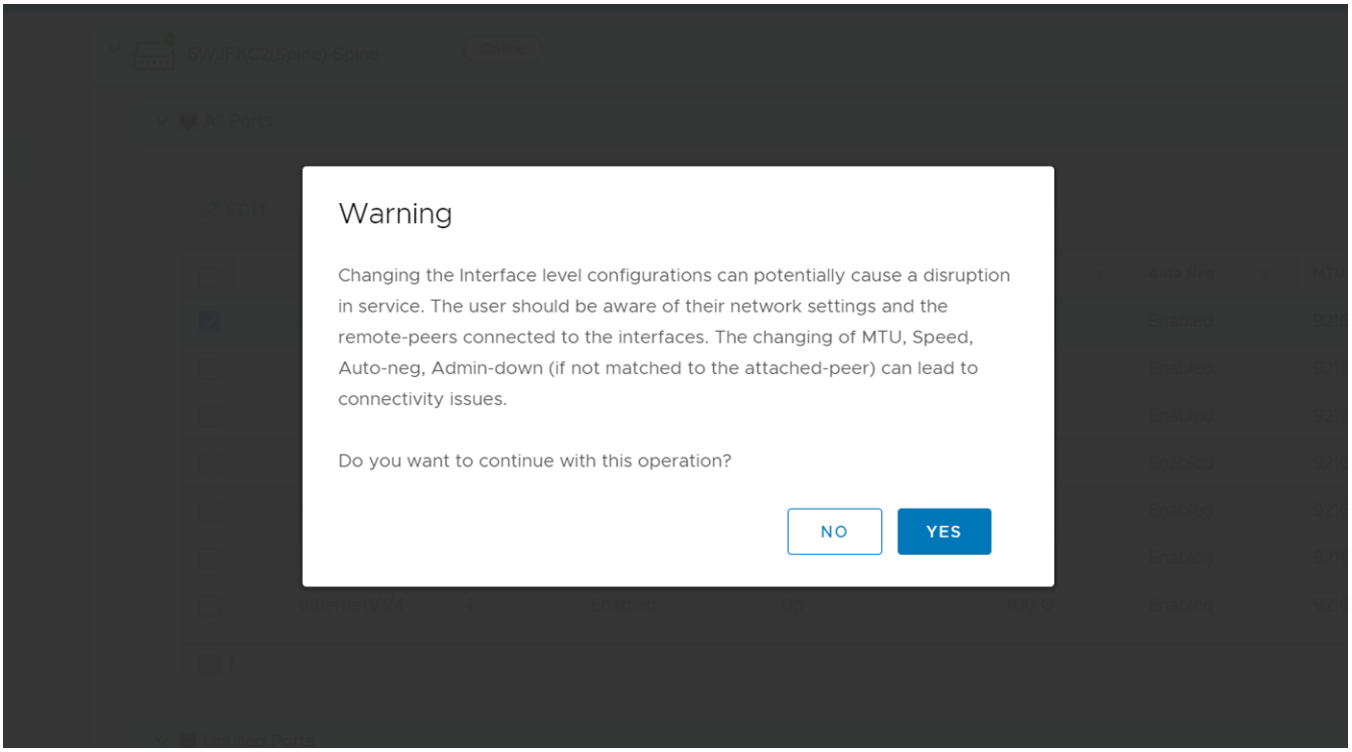

4. The system displays the action success or failure message.

Manage unused switch ports

You can view and manage the unused ports in the switches.

To enable or disable unused ports:

1. From the SmartFabric instance, select **Switches**.
2. Select any spine or leaf switch by clicking the arrow to view the list of ports.
3. Click **Unused Ports** category to view the list of unused ports available in the switch.
4. Select a port or multiple ports, and click **Enable Admin Status**.

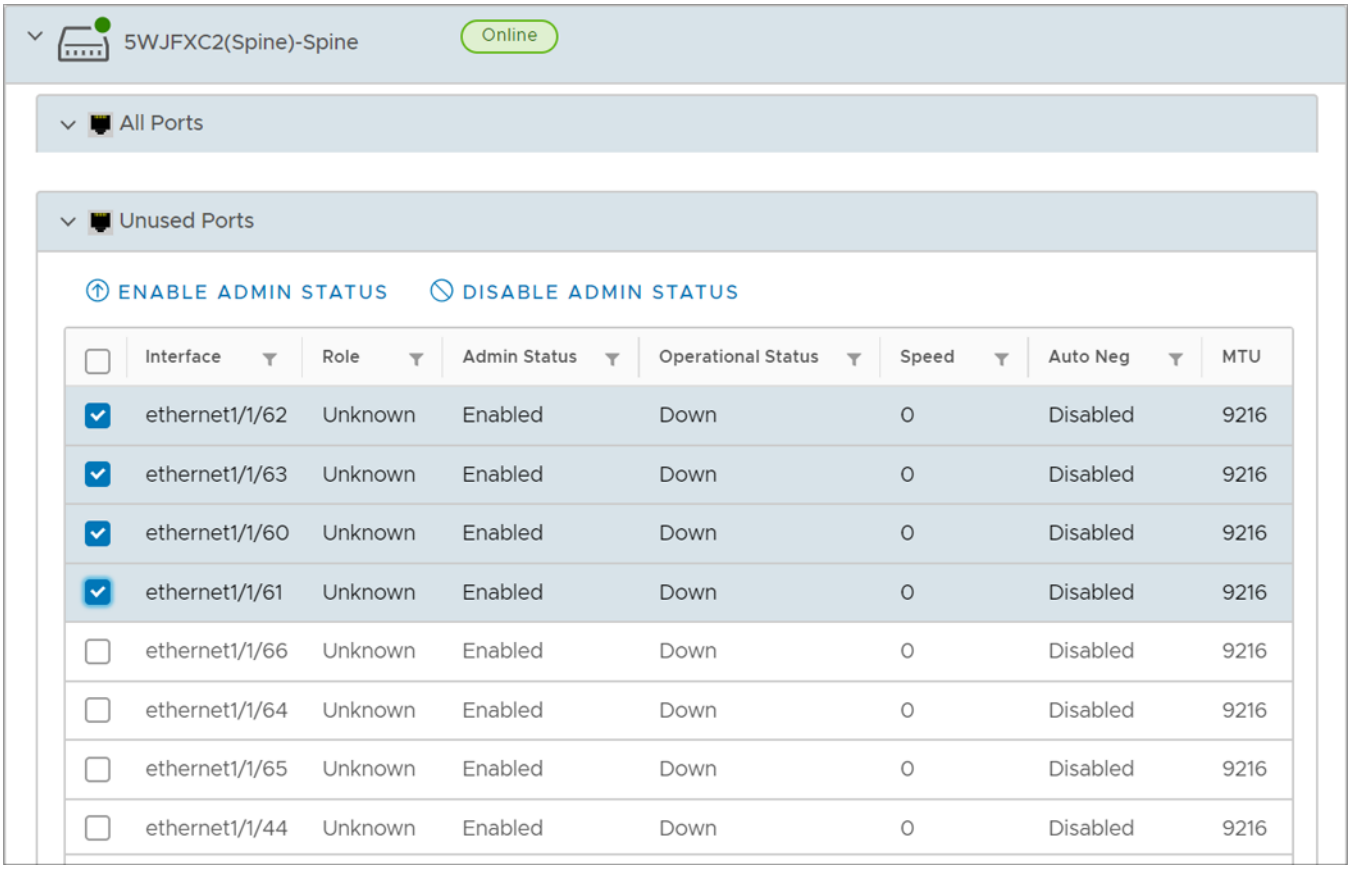

To disable the ports, select a port or multiple ports, and click **Disable Admin Status**.

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 OMNI UI.
  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.

Configure breakout ports

Configure breakout ports on an interface of the leaf switch.

NOTE: By default, the auto breakout feature is enabled in spine switches. OMNI UI does not provide an option to break out ports in spine switches.

To configure the breakout ports in a leaf switch:

1. From the SmartFabric instance, select **Switches** .
2. From **Leaf Switches** category, select a leaf switch from the list.

3. From **Breakout Port and Jump Port** category, select a port that you want to breakout, and click **Breakout port**.

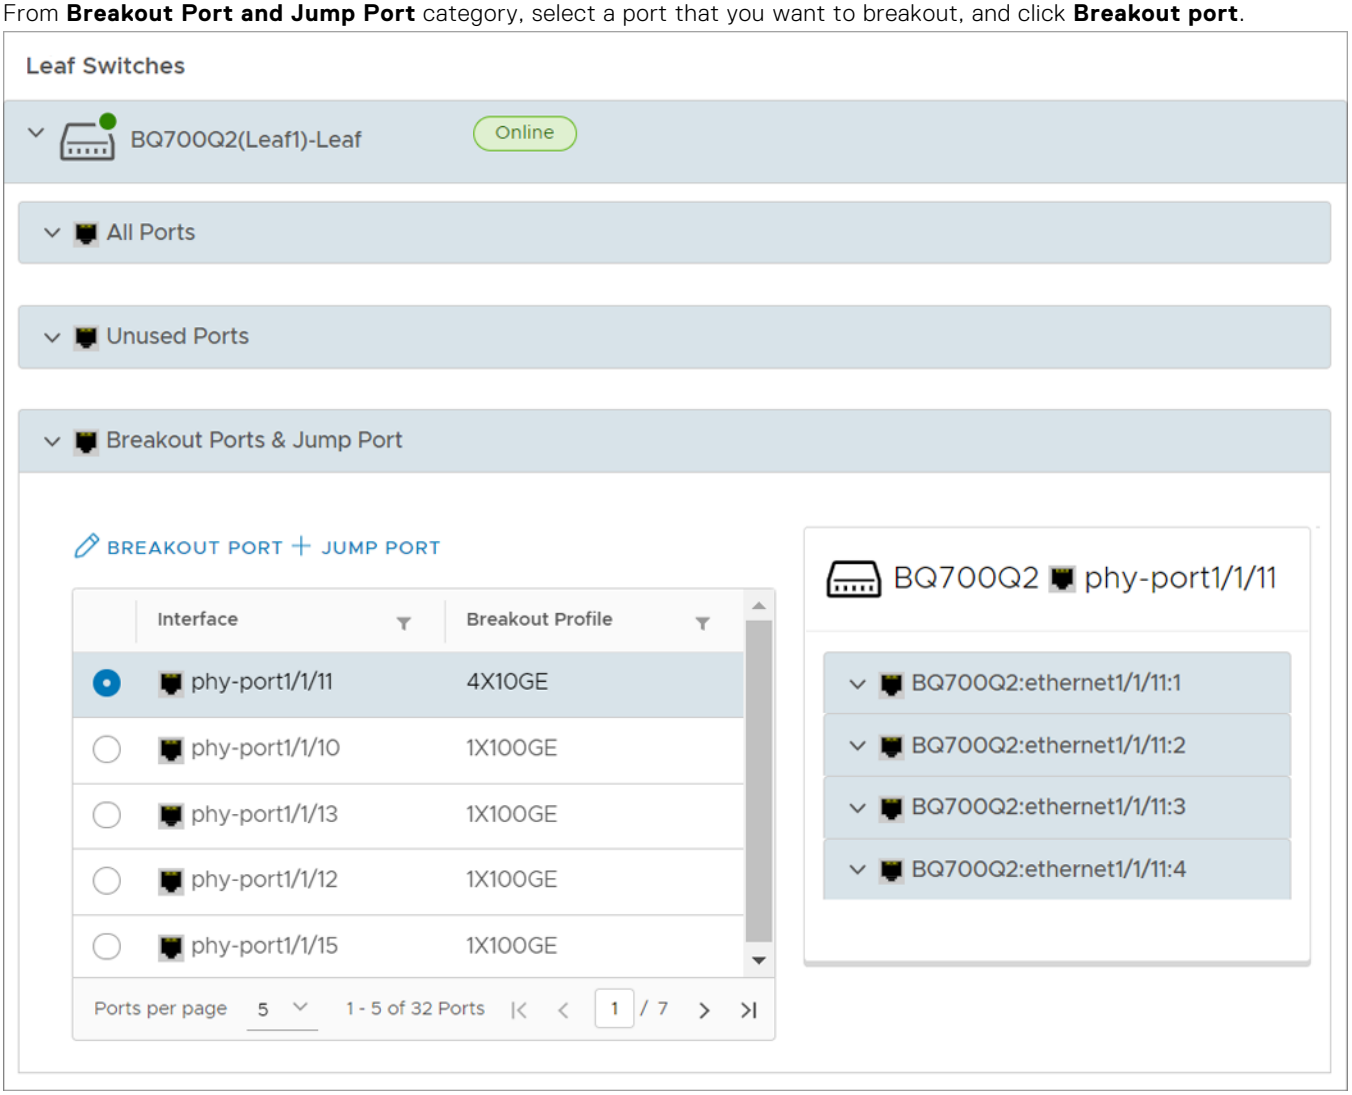

NOTE: While configuring a breakout port, the existing configuration of the port is reset to default.

4. Select the **Breakout Mode** for the port from the list, and click **Submit**.

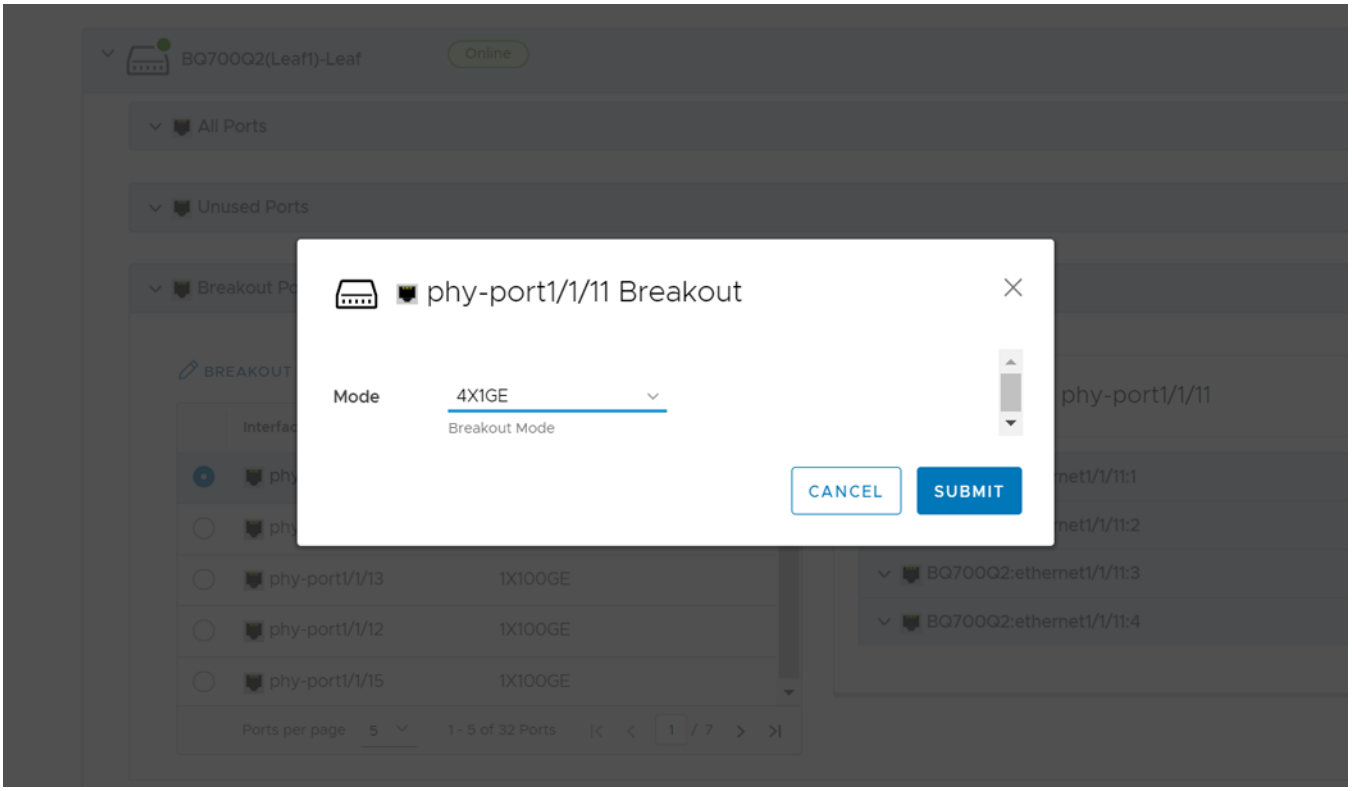

5. The system displays breakout port configured successful or failure message.

View port-group properties

Select a port to view properties on the right.

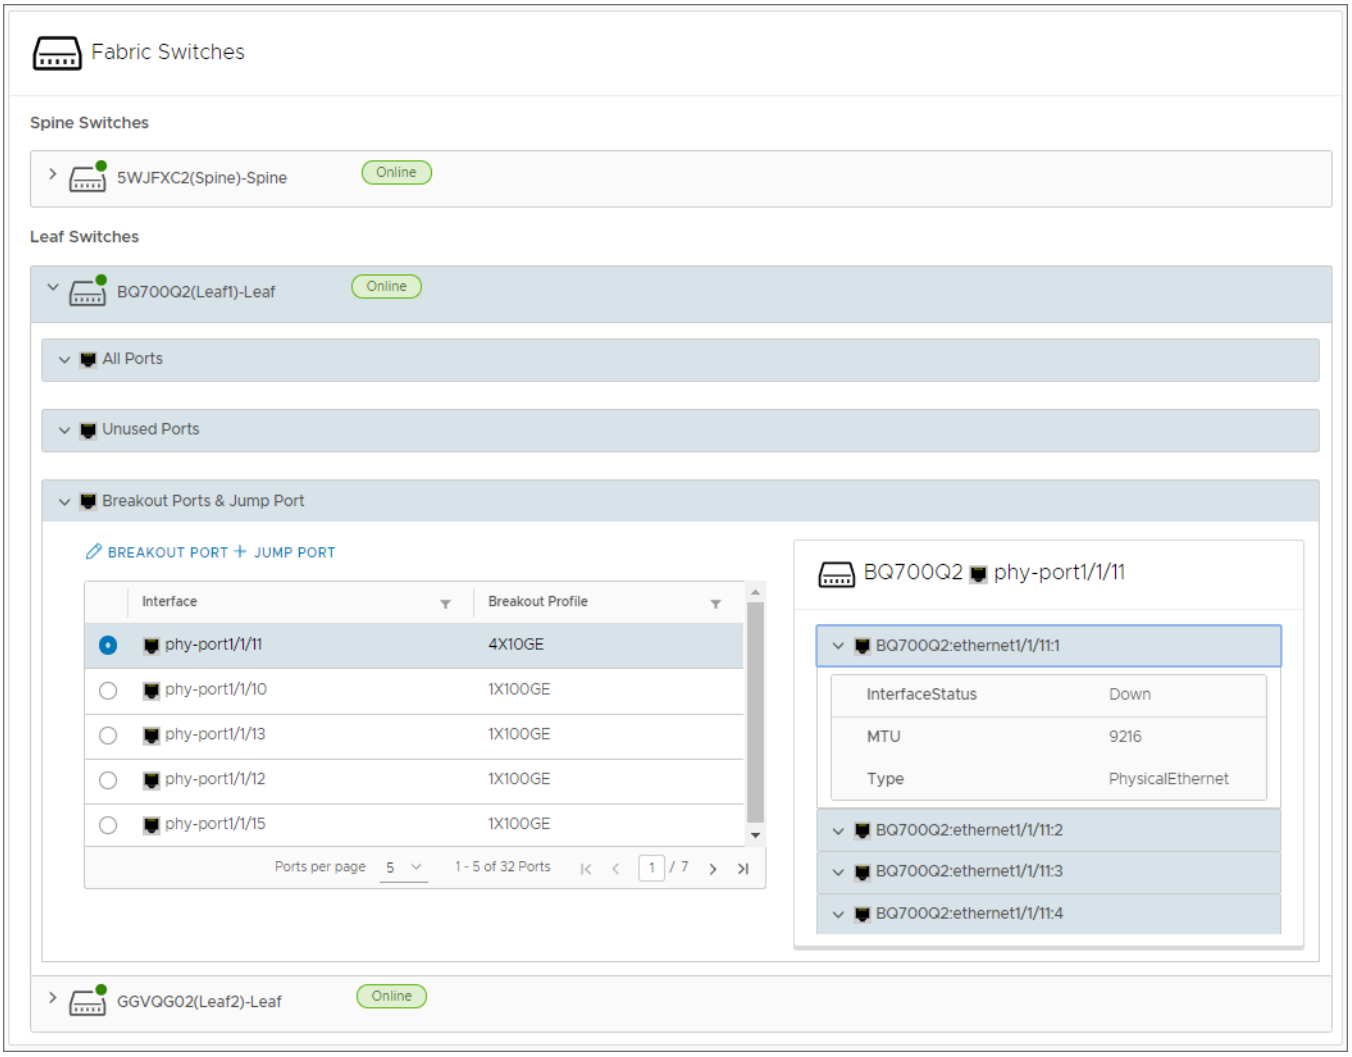

Add a jump port

You can configure only one port in a leaf switch as a jump port. You can select any available port that is not part of an uplink and ICL, and port connected to a server in SmartFabric deployment.

To configure a jump port:

1. Select the leaf switch from the list, and select the **Breakout Ports & Jump Port** category.

2. Select the switch to view the properties, and click **Jump Port**.

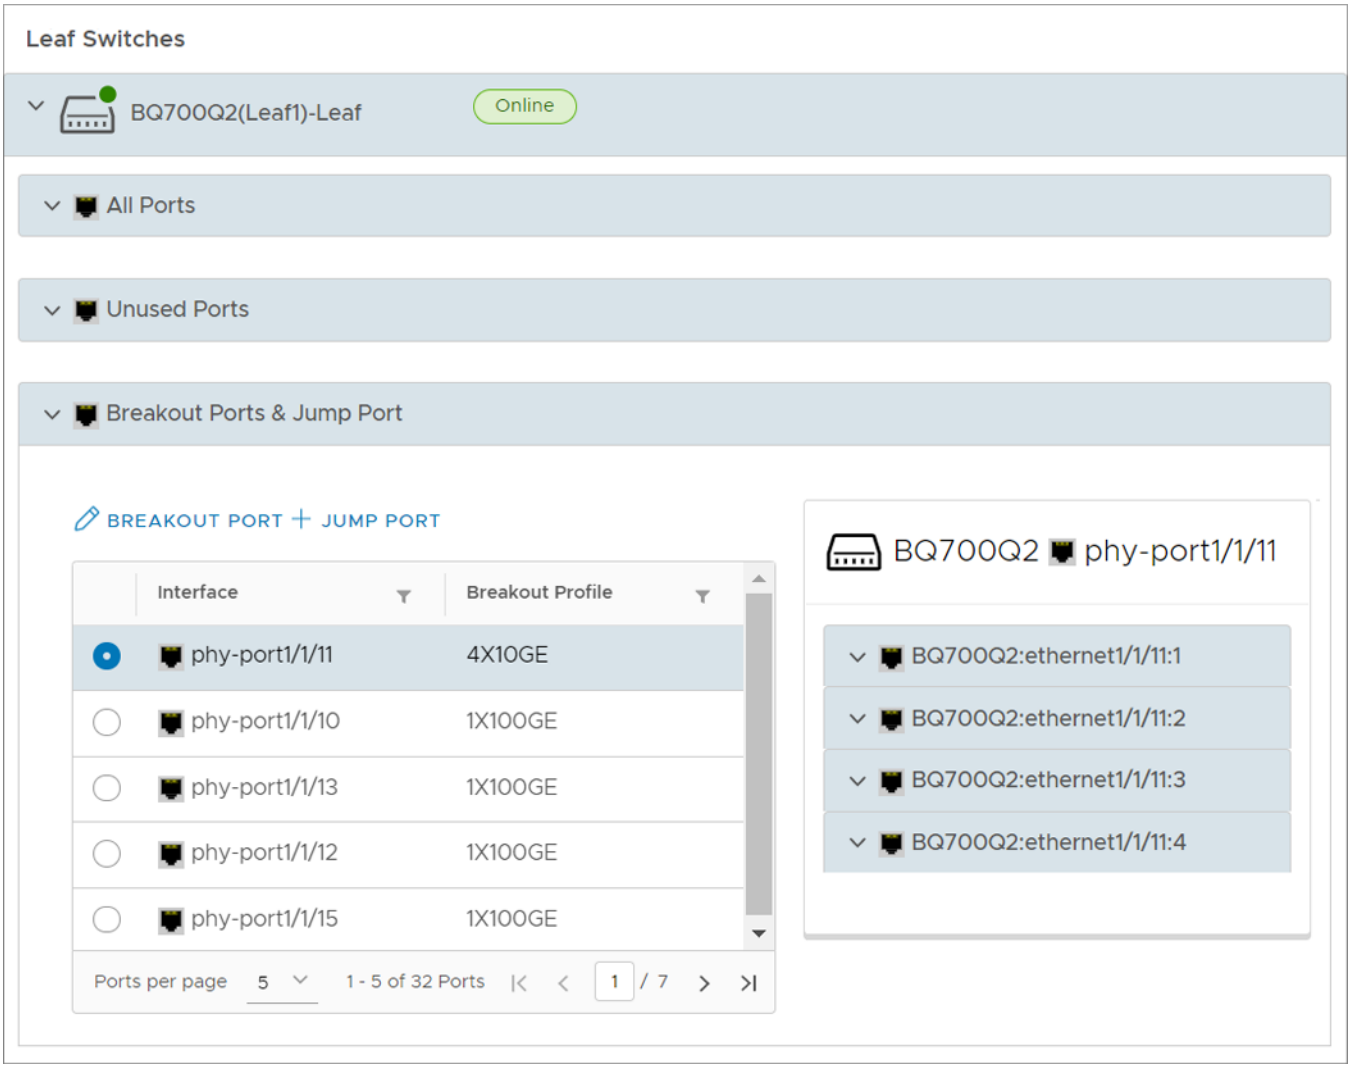

3. Enter the **Name** of the new jump port, select the **Interface Name**, select the **Untagged Network**, then click **Add**.

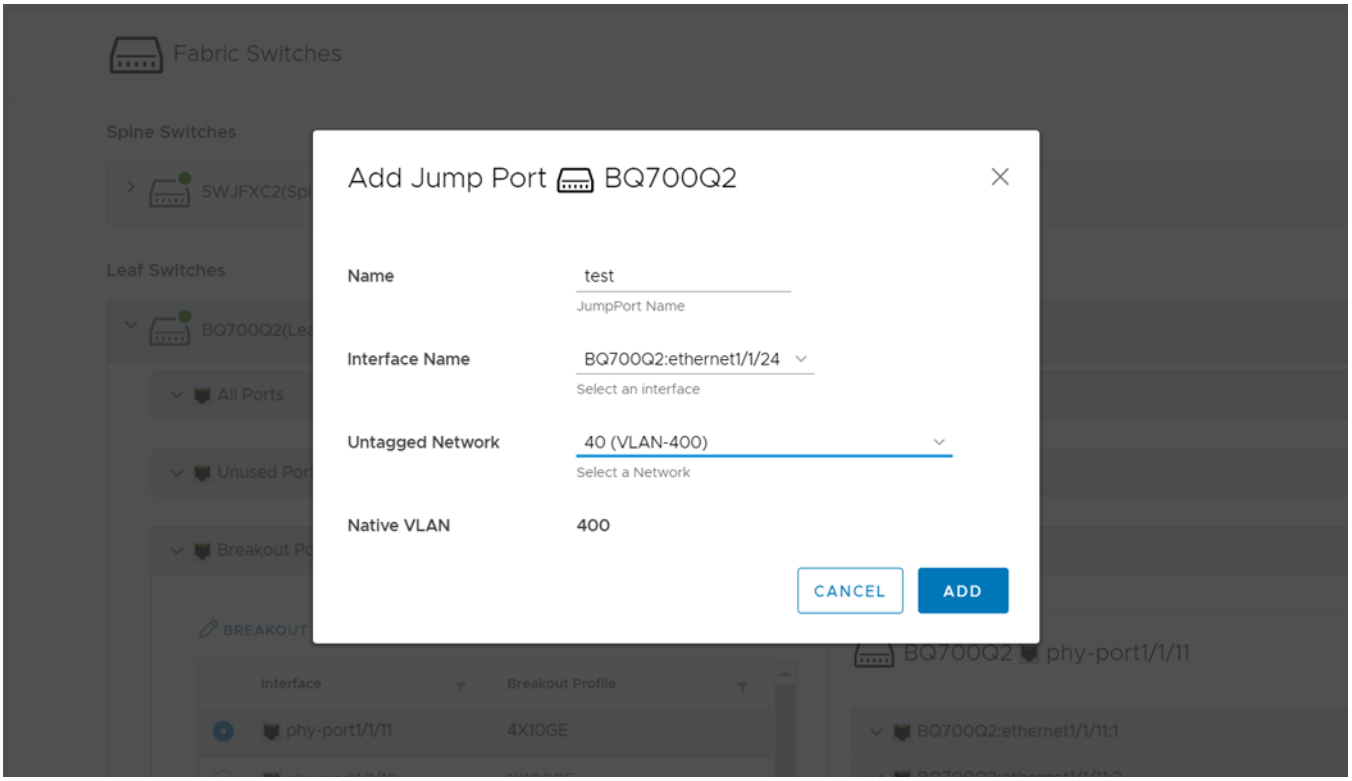

4. The system displays jump port addition success message.

Delete jump port

1. Select the leaf switch for which you want to delete the configured jump port.

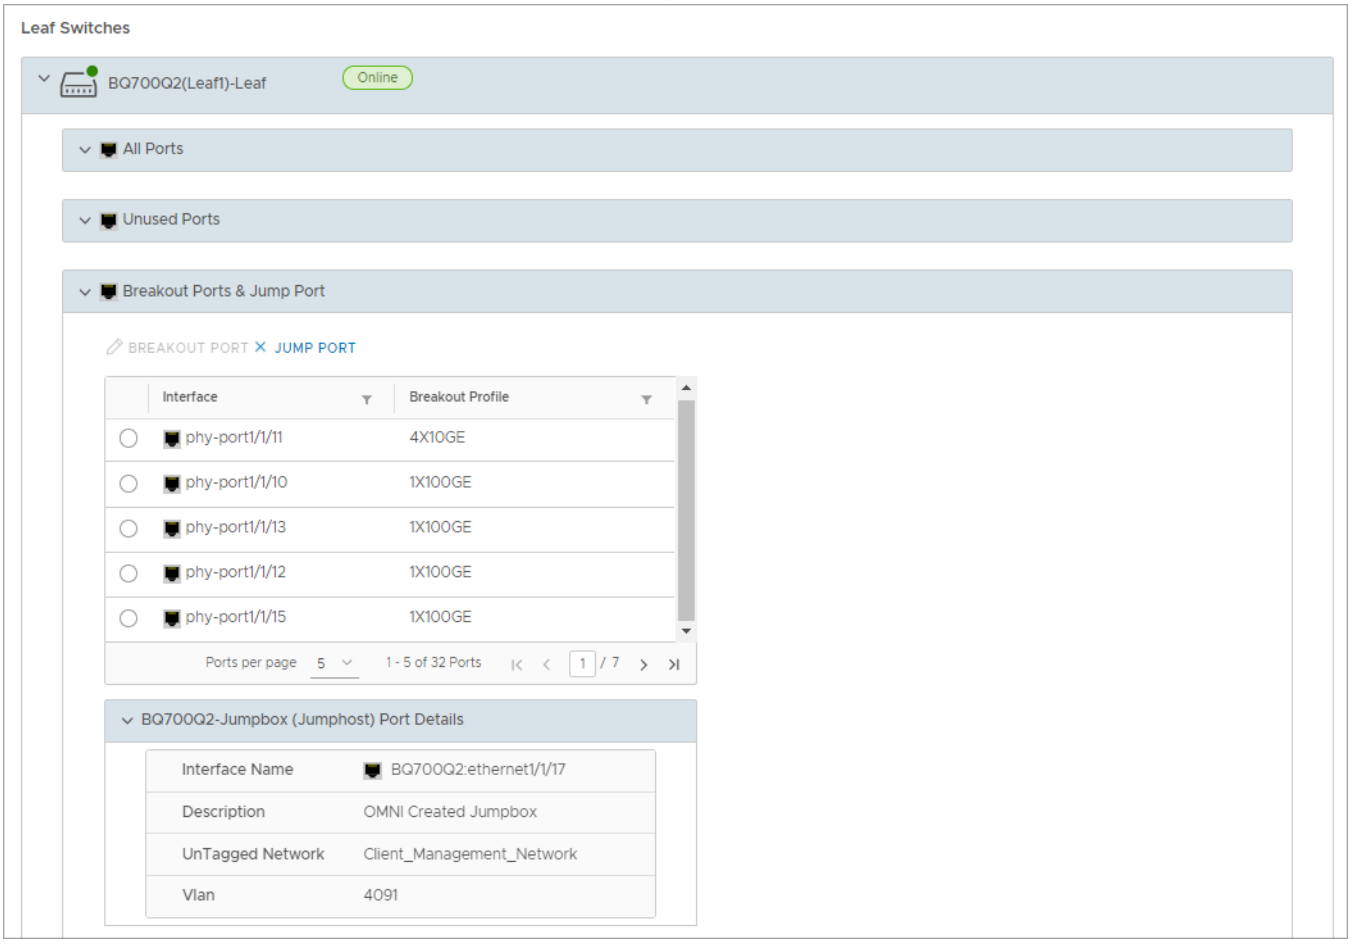

2. Select the Jump port, and click **Delete**.

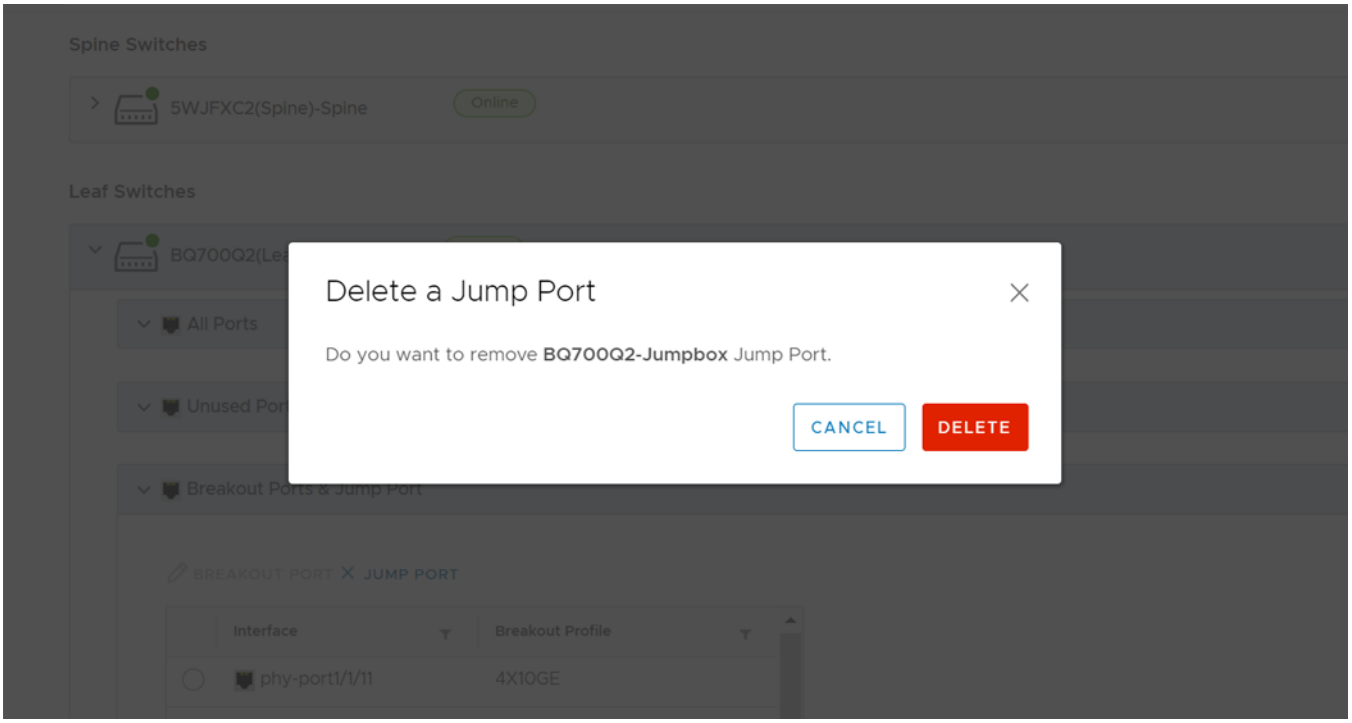

3. The system displays jump port deletion success message.

Configure server interface profile

Server Interfaces Profile page displays a list of Server Profile IDs and their respective onboard status. Select a profile to view details pertaining to that specific profile. You can view information including interface ID, fabric ID, native VLAN, and network name and VLAN ID (if applicable).

From **Server Interface**, you can:

- Create a server interface profile.
- Edit a network in a server interface profile.
- Edit the ports in a server interface profile.
- Delete a server interface profile.
- Automate server onboarding.

Create server interface profile

Create a server profile by providing the server profile type, name, and bonding technology.

Create server interface with an existing server profile

To create a server interface with an existing server profile:

1. From SmartFabric instance, select **Server Interface**.

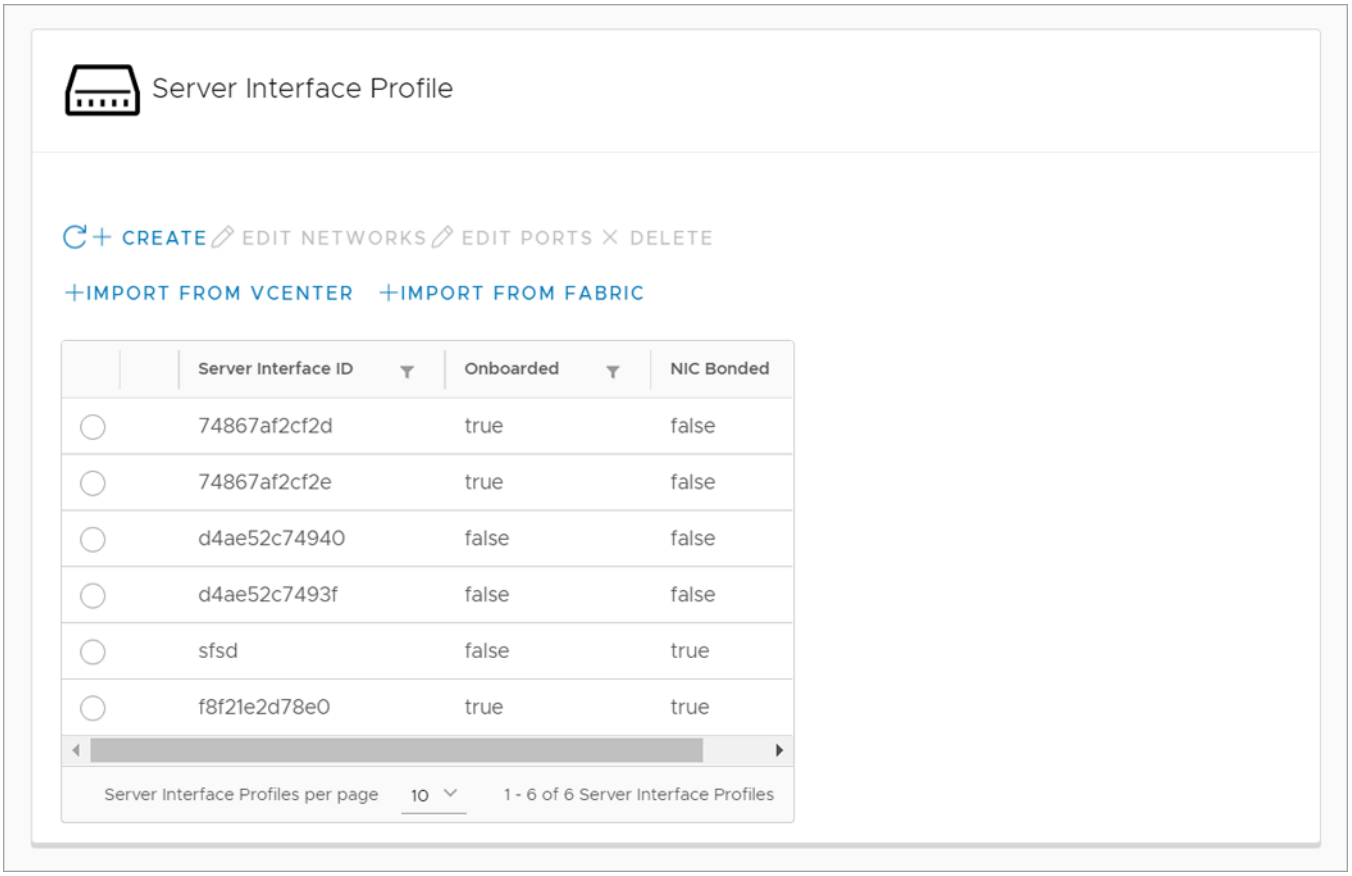

2. Click **Create** to create a server interface profile and provide server interface ID, then select **Existing Server Profile**.

   NOTE: You cannot configure duplicate server interface ID. When using MAC address to onboard server interface, enter MAC Address without ":", for example, f8f21e2d78e0. For onboarding ESXi host Interfaces for zero touch automation, use the ESXi host VM NIC physical adapter MAC address without ":".

3. Select the **Server Profile Id** from the list, select one or multiple networks for the **Untagged Network**, enable or disable **NIC Bonding**, select **Static Onboarding Option** as **No**, and click **Create**.

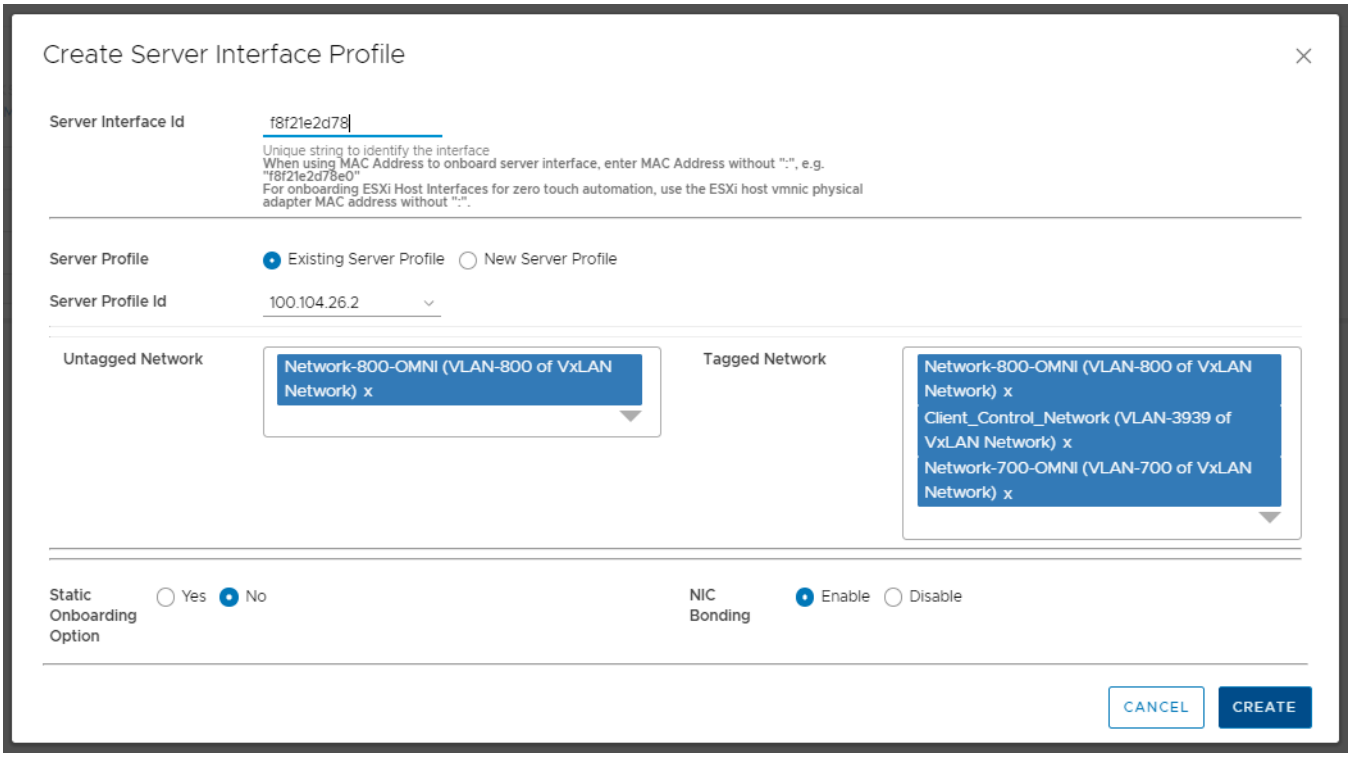

4. (Optional) Select **Yes** for the **Static Onboarding Option**, add **Leaf Node** and **Interface** (where the server interface is connected), select the routing protocol as **None**, and click **Create**.

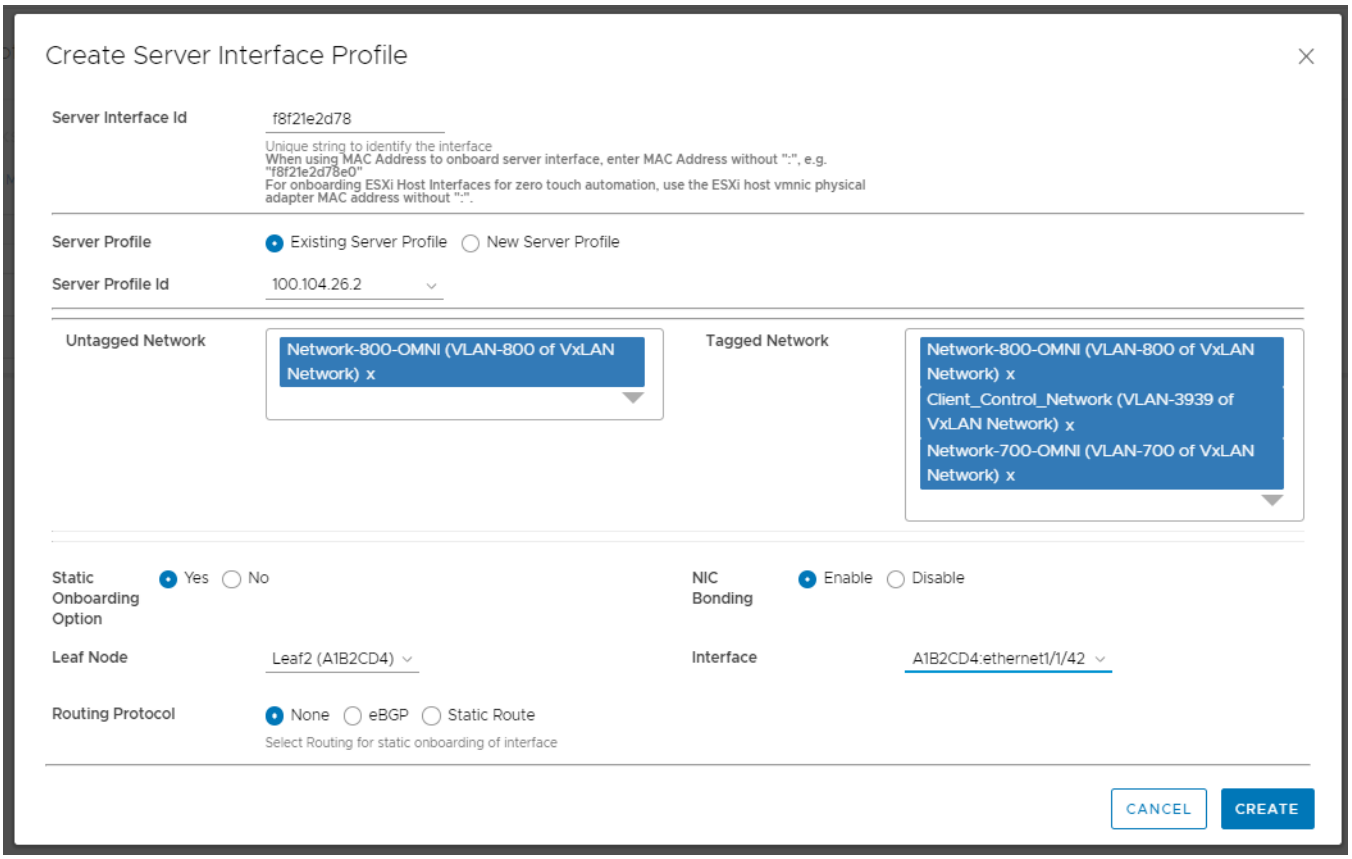

5. (Optional) Select **Yes** for the **Static Onboarding Option**, select **Leaf Node** and **Interface** (where the server interface is connected), select **eBGP**. Enter the eBGP routing template by entering the name, peer **ASN**, description, and peer interface **IP address**, and click **Create**.

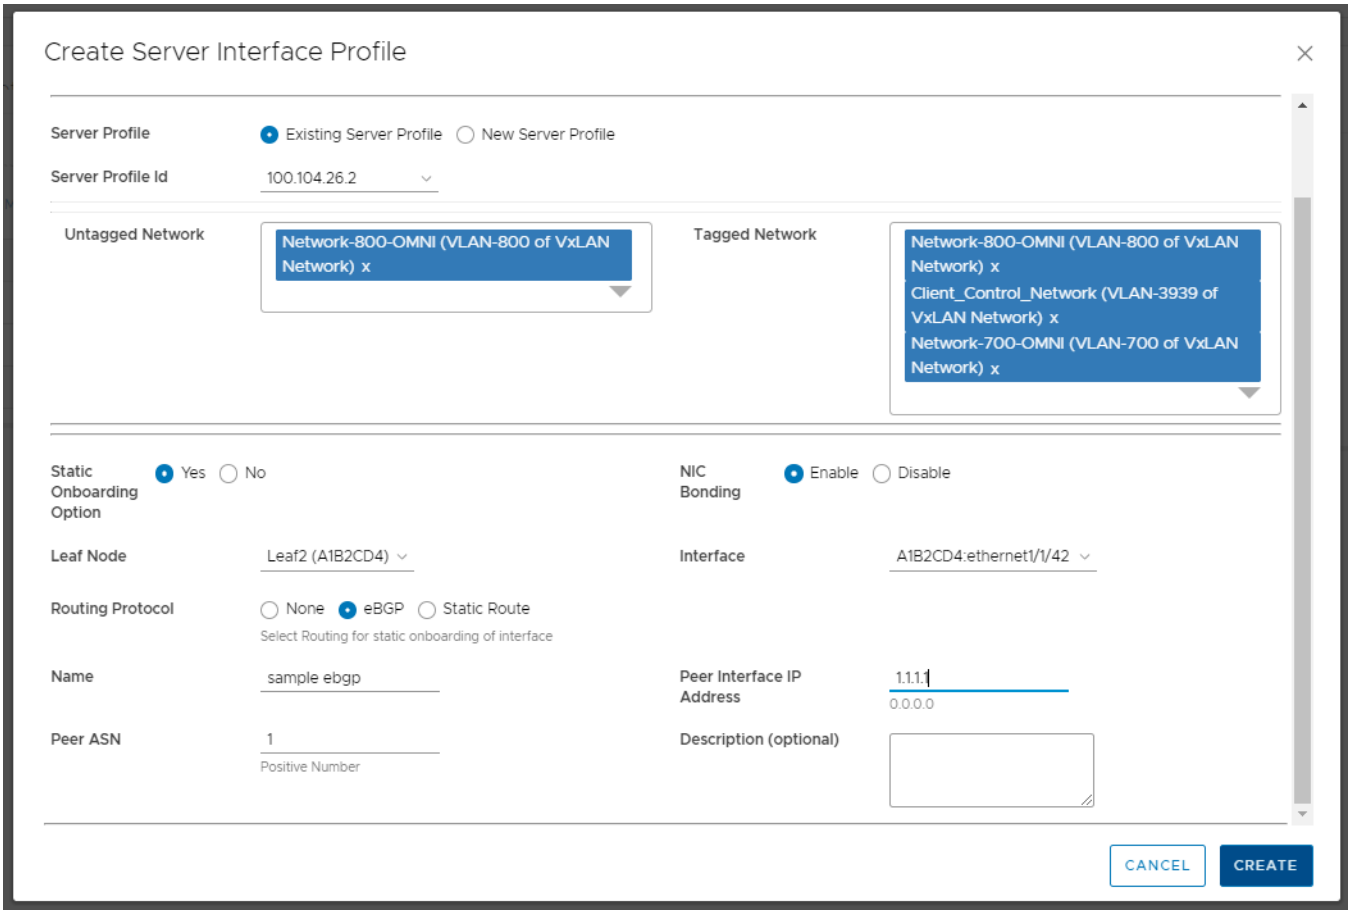

NOTE: In static onboarding, the eBGP or static route routing protocol option is used for NSX-T deployment.

6. (Optional) Select **Yes** for the **Static Onboarding Option**, select **Leaf Node** and **Interface** (where the server interface is connected), select **Static Route**, enter the **Network Address** and **Next-Hop Address**, then click **Create**.

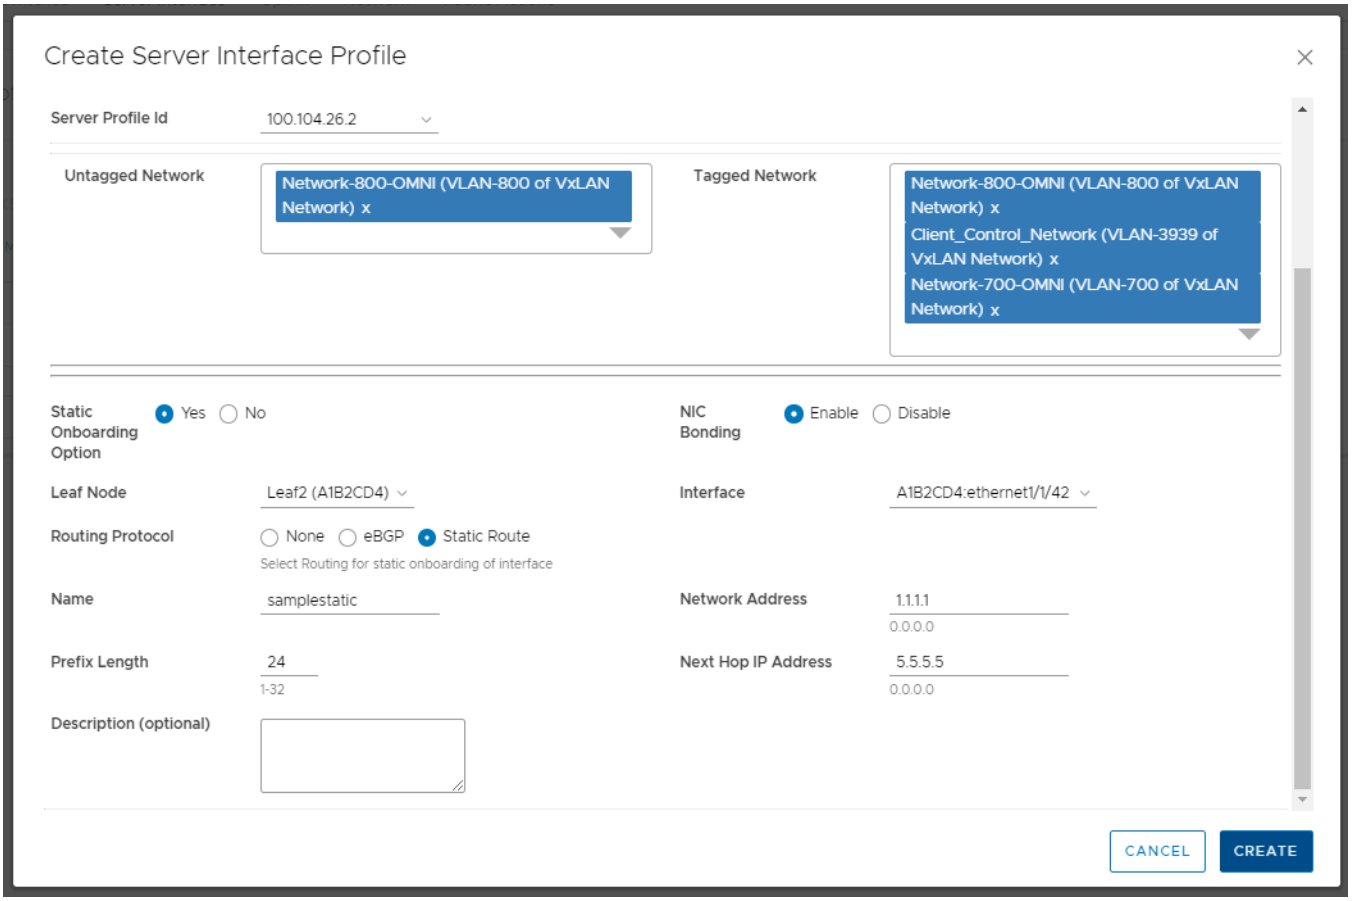

NOTE: You cannot delete any created server profile.

7. The system displays server profile and server interface creation successful messages.

Create server interface with new server profile

To create a server interface with new server profile:

1. From SmartFabric instance, select **Server Interface**.
2. Click **Create** to create a server interface profile and provide server interface ID, then select **New Server Profile**.

NOTE: You cannot configure duplicate server interface ID. When using MAC address to onboard server interface, enter MAC Address without ":", for example, f8f21e2d78e0. For onboarding ESXi host Interfaces for zero touch automation, use the ESXi host VM NIC physical adapter MAC address without ":".

3. Select the **Server Profile Id** and **Server Profile Bonding Type** from the list, select the **Untagged Network** and **Tagged network**, enable or disable **NIC Bonding**, select **Static Onboarding Option** as **No**, and click **Create**.

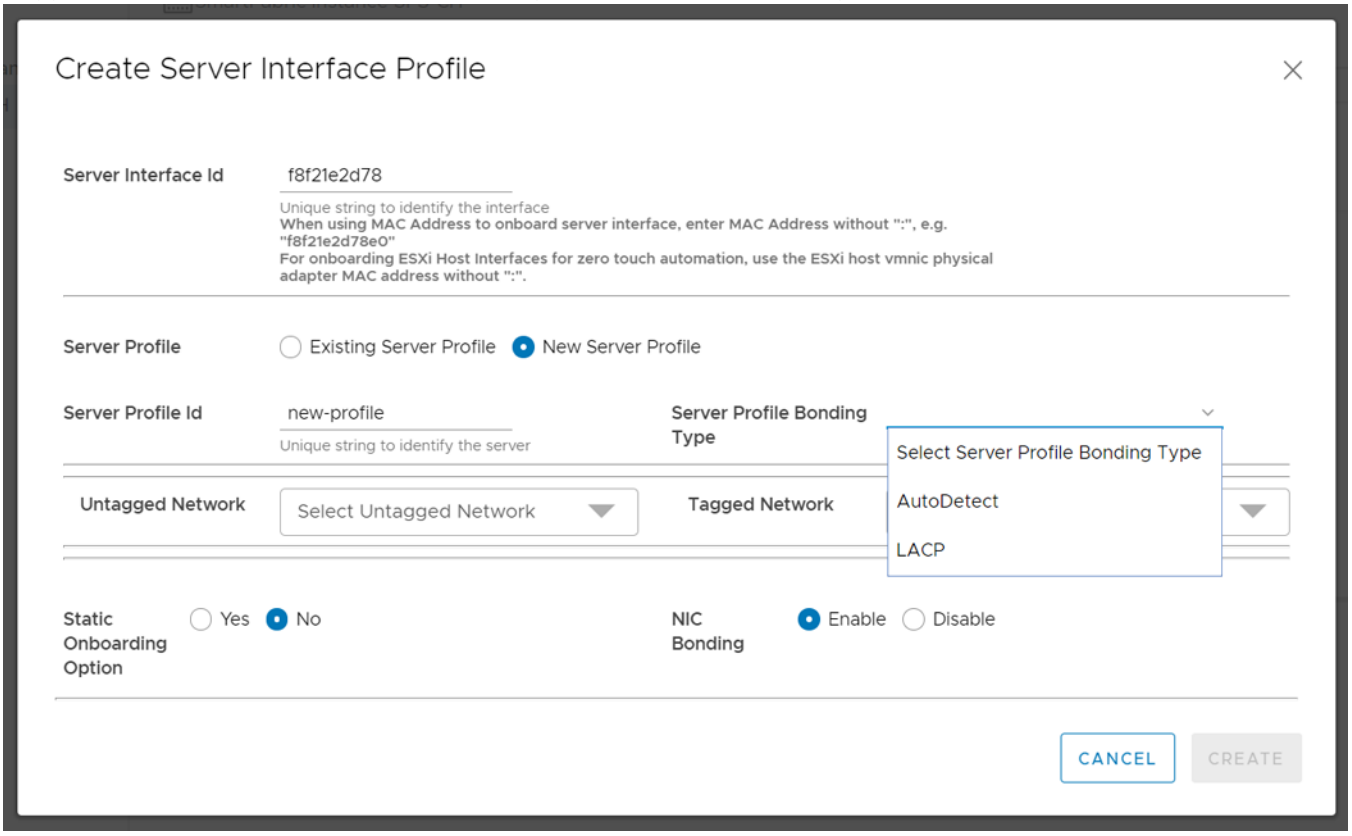

4. (Optional) Select **Yes** for the **Static Onboarding Option**, add **Leaf Node** and **Interface** (where the server interface is connected), select the routing protocol as **None**, and click **Create**.

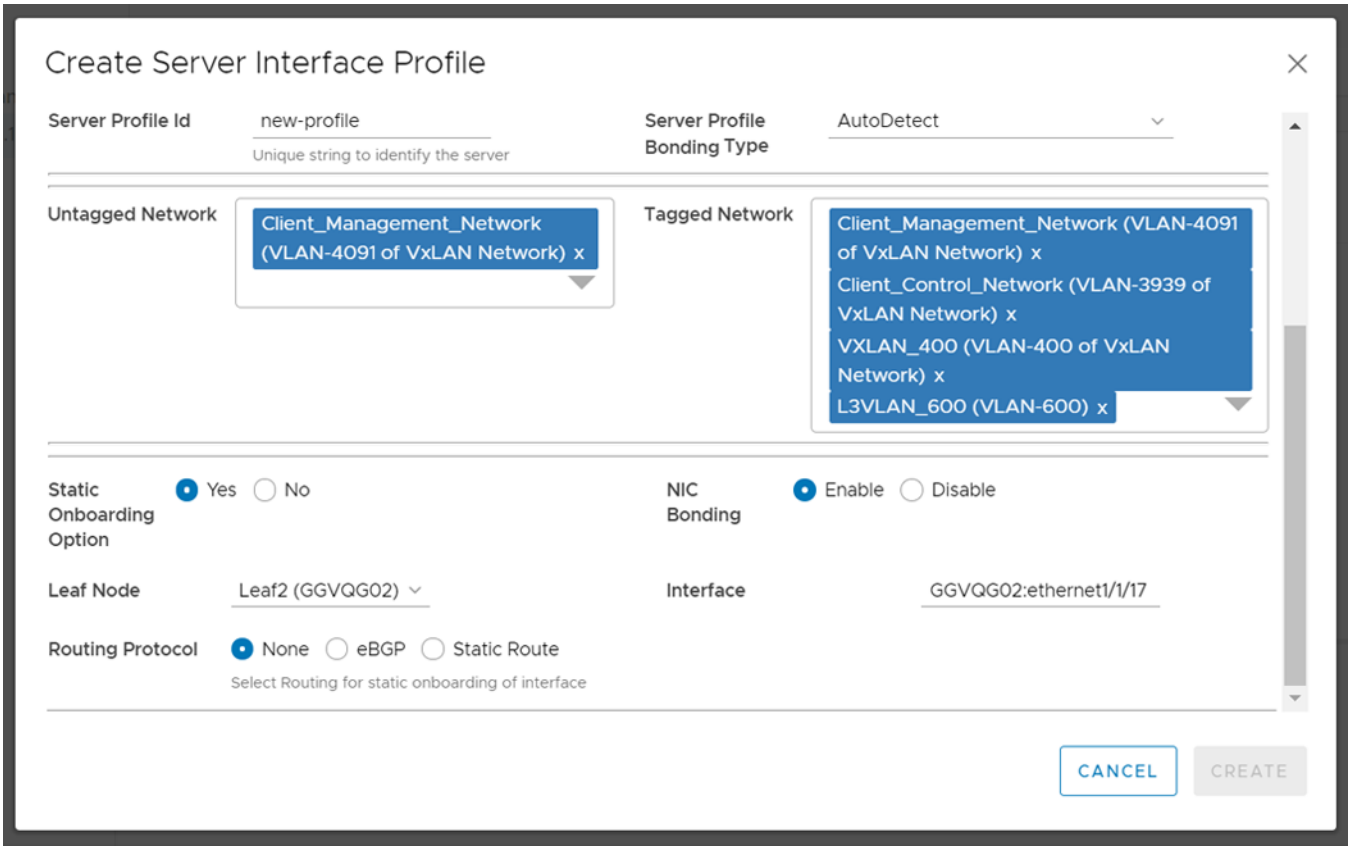

5. (Optional) Select **Yes** for the **Static Onboarding Option**, select **Leaf Node** and **Interface** (where the server interface is connected), select **eBGP**. Enter the eBGP routing template by

entering the name, peer **ASN**, description, and peer interface **IP address**, and click **Create**.

NOTE: In static onboarding, the eBGP or static route routing protocol option is used for NSX-T deployment.

6. (Optional) Select **Yes** for the **Static Onboarding Option**, select **Leaf Node** and **Interface** (where the server interface is connected), select **Static Route**, enter the **Network Address** and **Next-Hop Address**, then click **Create**.

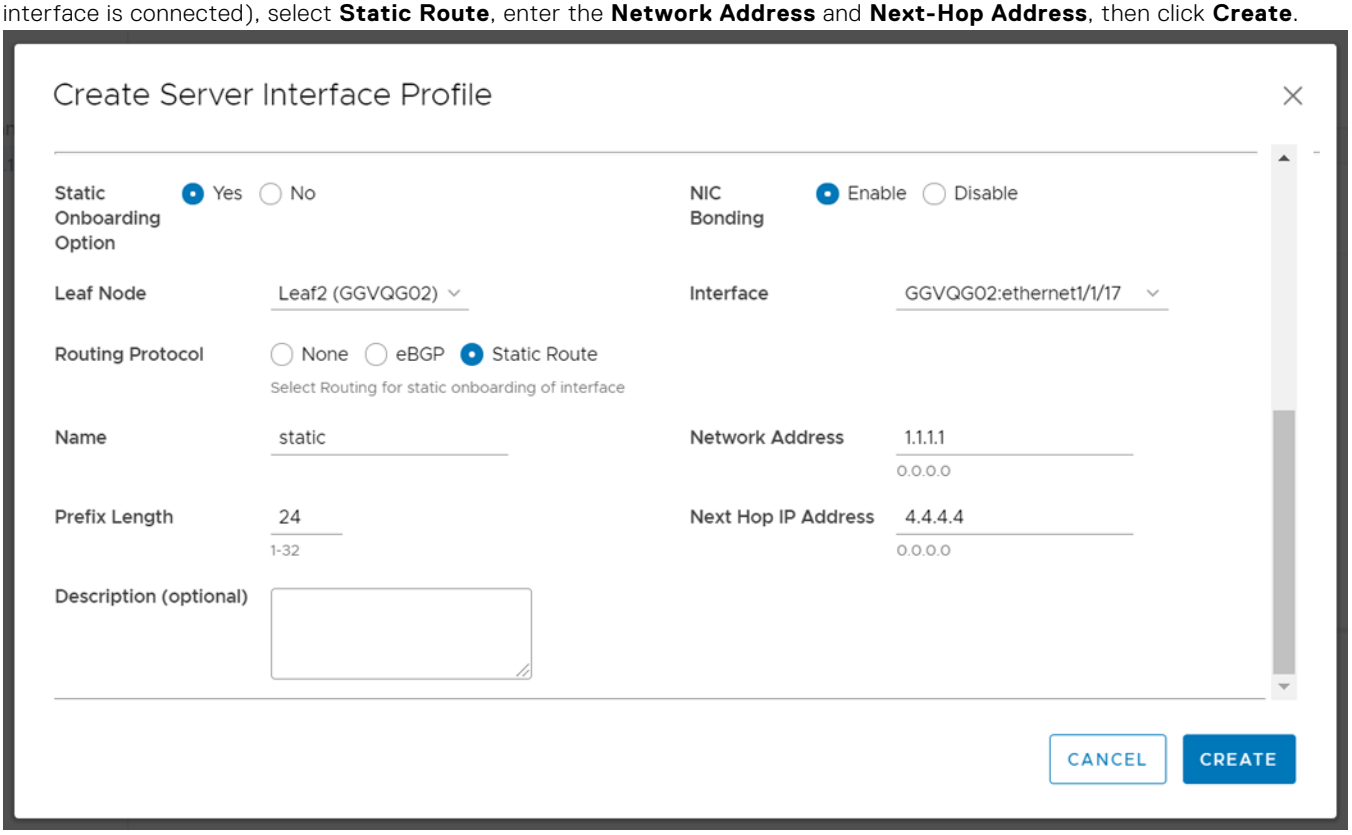

NOTE: You cannot delete any created server profile.

7. The system displays server profile and service interface creation successful messages.

NOTE: OMNI does not synchronize a statically onboarded interface when it is first added. For the synchronization to happen, a port-group change event on the vCenter must happen or a restart of the automation service for the specific vCenter and SmartFabric instance must occur.

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. From SmartFabric instance, select **Server Interface**.
2. Select the server interface ID from the list to view the detailed information.

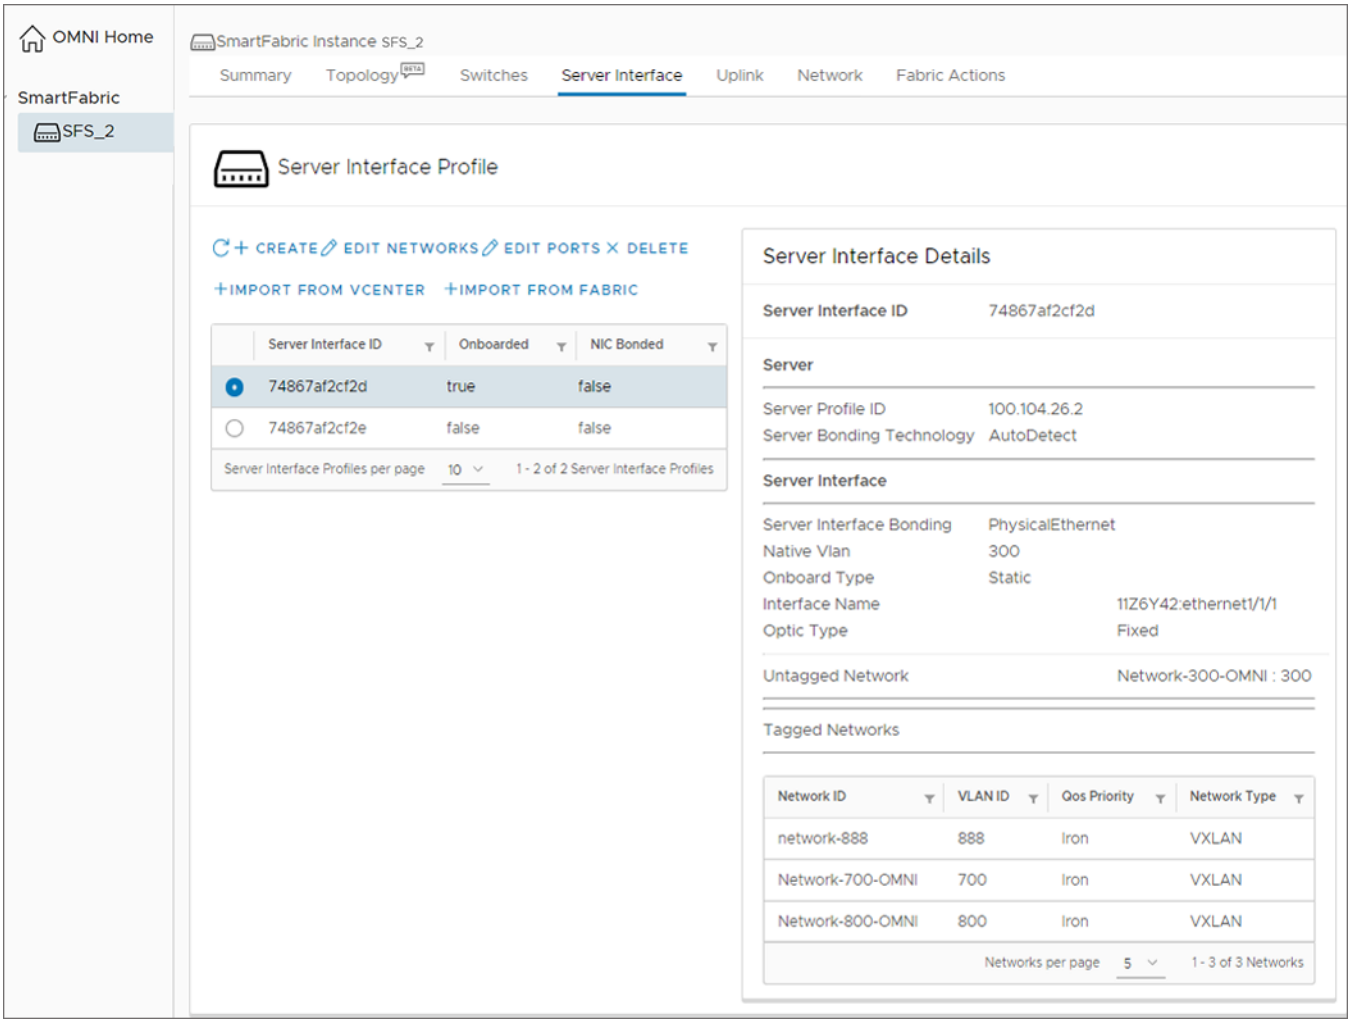

3. Select the server interface ID from the list, and click **Edit Networks**.

4. Edit the **Untagged Network** and the **Network** configuration for the profile, and click **Edit**.

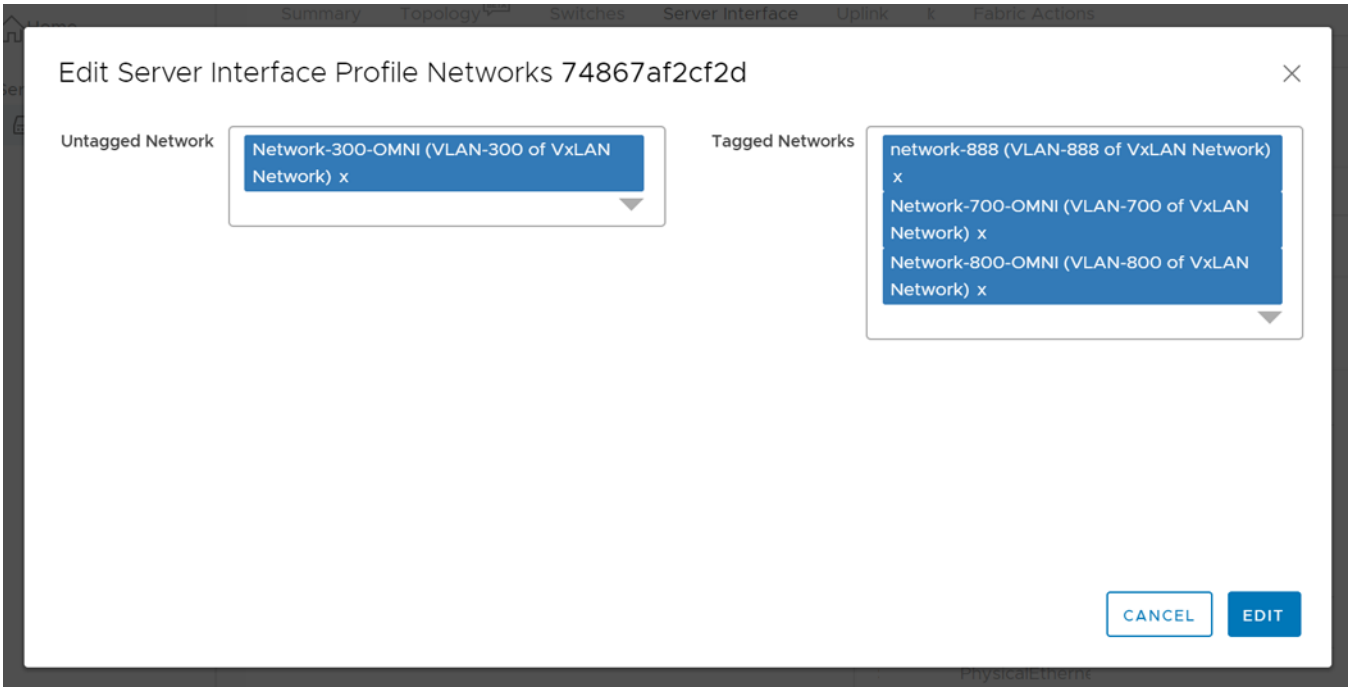

5. The system displays the server interface profile update success message.

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, and click **Edit**.

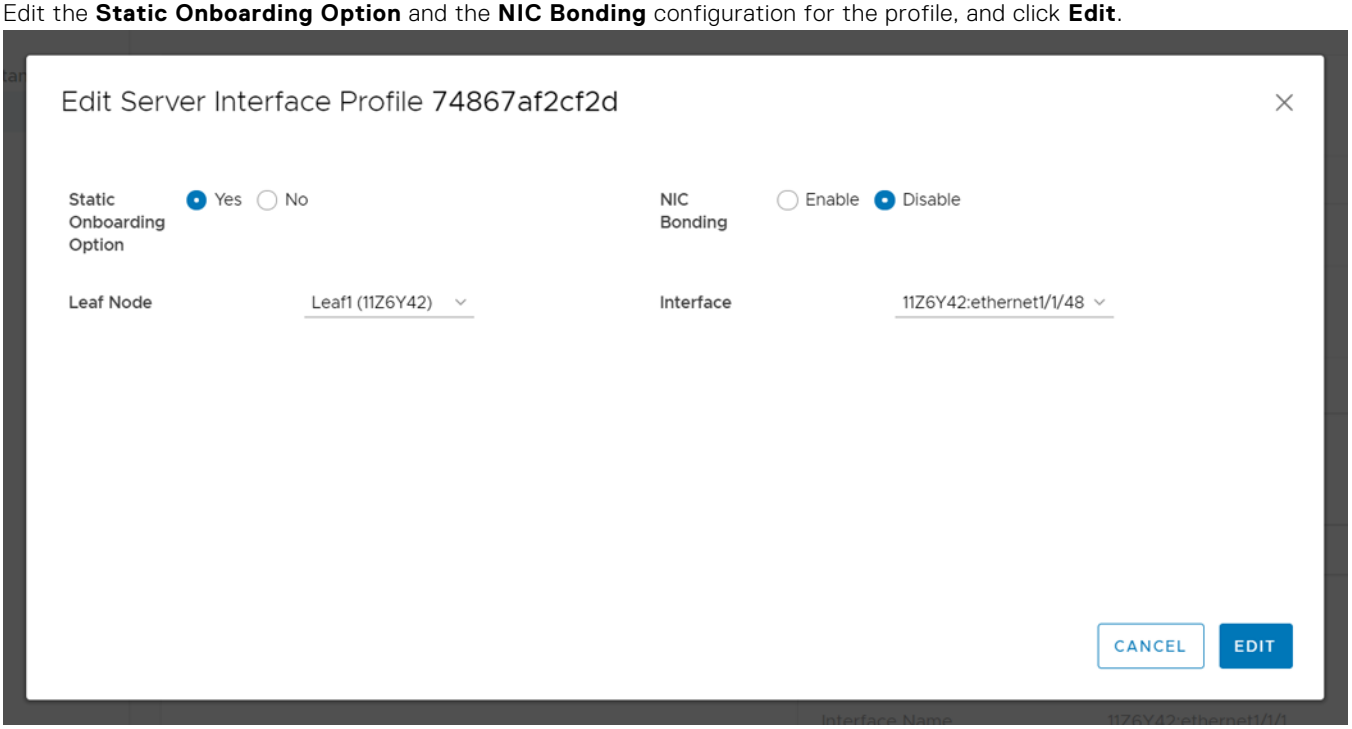

3. The system displays the server interface profile update success message.

Delete a server interface profile

You can delete a service interface profile from the SmartFabric instance. To delete:

1. Select the server interface profile from the displayed list, and click **Delete**.

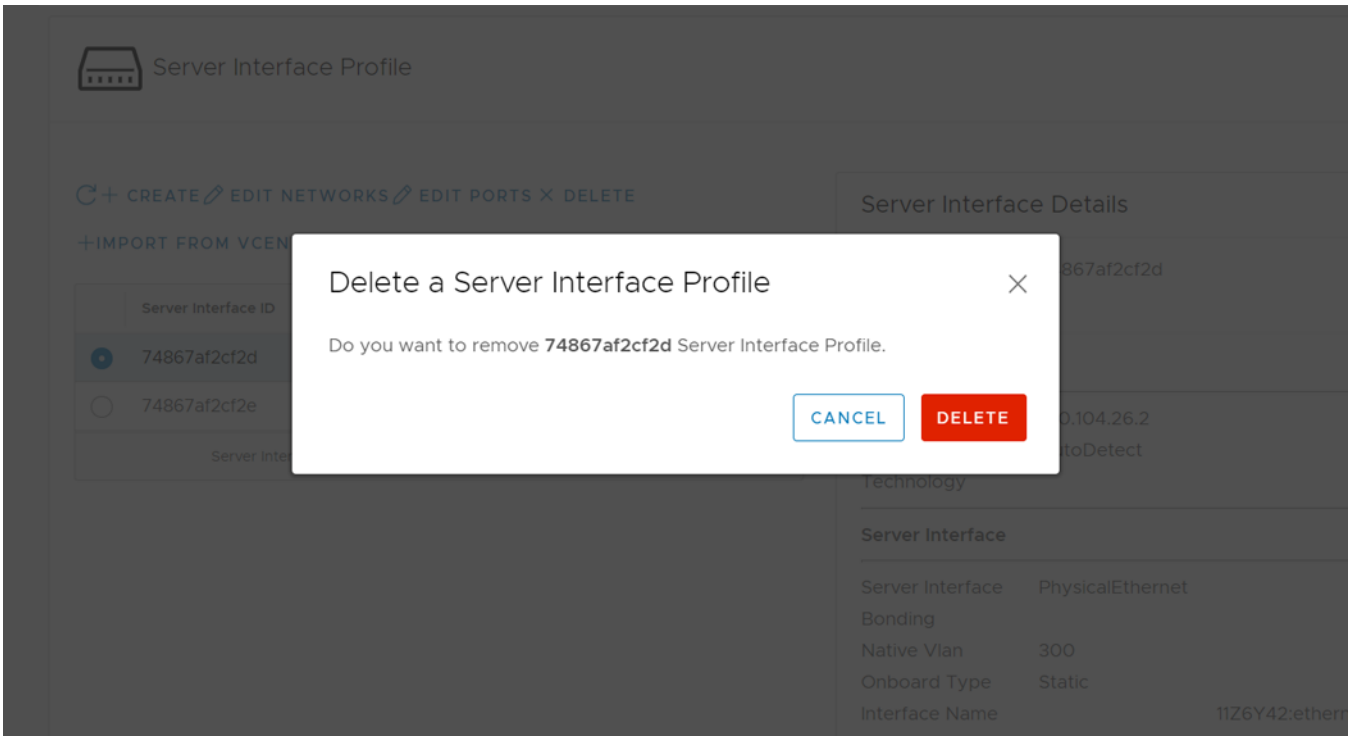

2. Click **Delete** to confirm.

Import ESXi host profiles from vCenter

Automate onboarding of server interface profile by importing:

Use this feature to migrate the existing ESXi hosts that are already connected to the vCenter and ready to be onboarded on to the fabric. The feature imports all the required servers to onboard on to the SFS instead of manually configuring the server interface one at a time.

OMNI retrieves data center, clusters, hosts, VM NICs, and networks for the registered vCenter. Create server interface profiles for the set of available VM NICs in ESXi hosts from vCenter.

NOTE: In vCenter, enable LLDP on Distributed Virtual Switch of ESXi host to discover the interfaces automatically.

1. From SmartFabric instance, select **Server Interface**.

2. Click **Import from vCenter** to launch the **Onboarding ESXi Hosts** wizard.

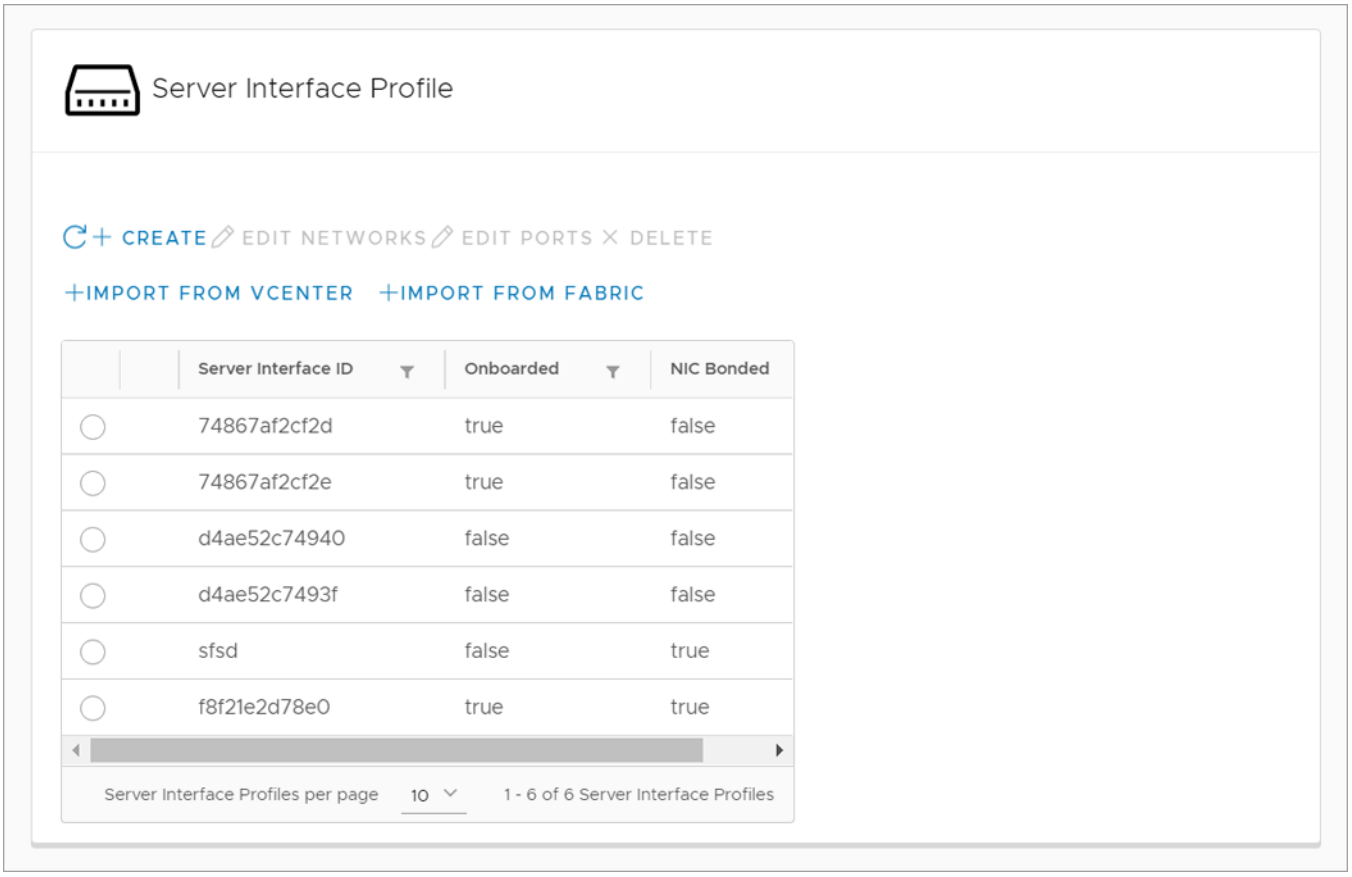

3. Select the **vCenter** from the list, and click **Next**.

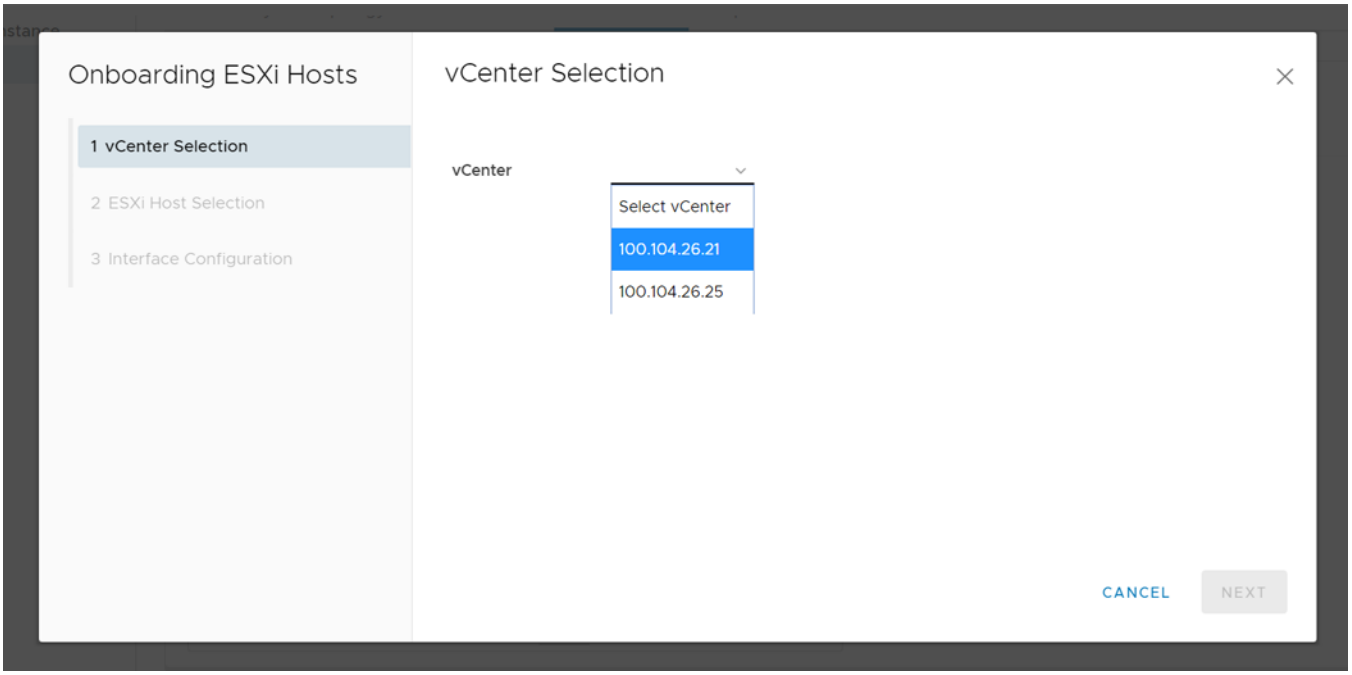

4. Select the relevant cluster, the ESXi host, or the VM NICs available on the ESXi host. **ESXi Host Selection** window displays the server profile status of the interfaces on the right.

NOTE: You cannot select the VM NICs that are already part of a server interface profile in SmartFabric.

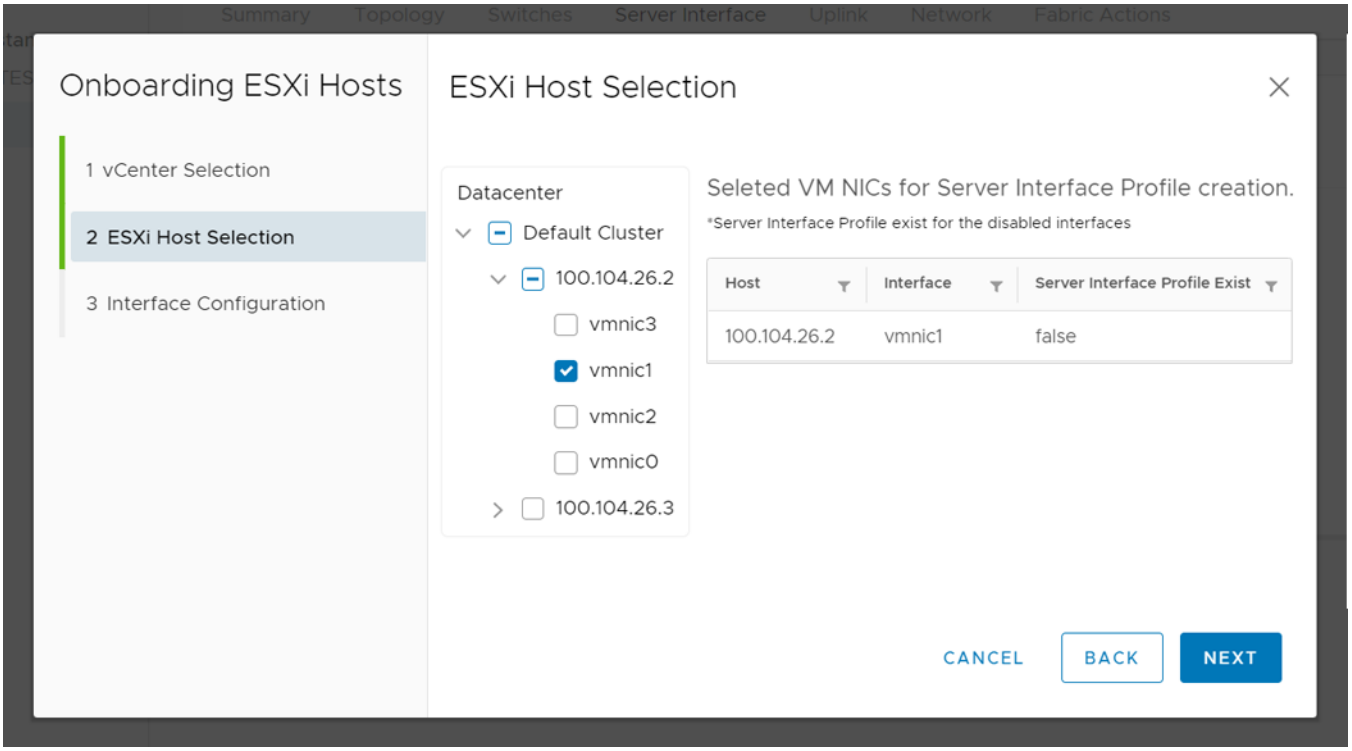

5. Click **Next** to complete the selection of the VM NICs.
6. The **Interface Configuration** screen displays the list of selected VM NICs.

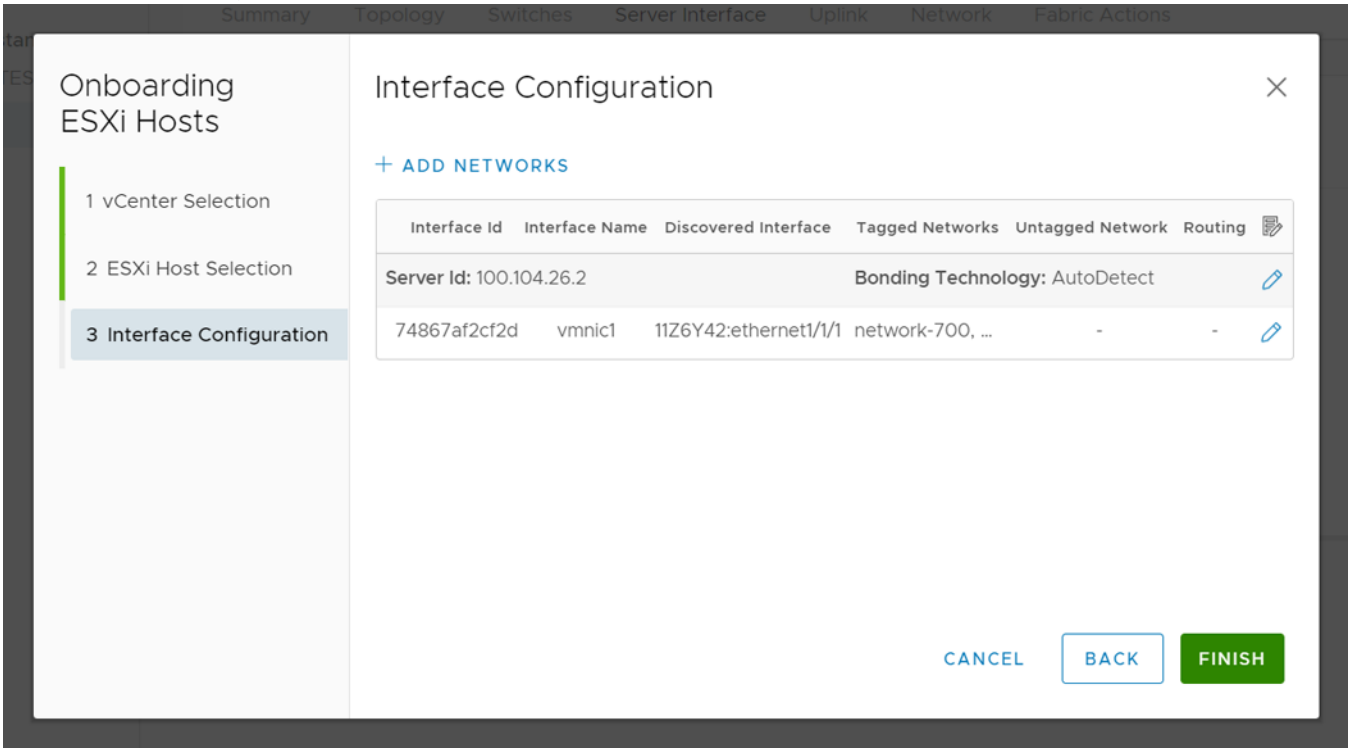

7. (Optional) Click **Edit** icon available for each interface to edit the server profile information.

   Edit the NIC bonding configuration and **Static Onboarding**. If the static onboarding is **No**, select an **Untagged Network** and one or more **Tagged Networks** and click **Update**.

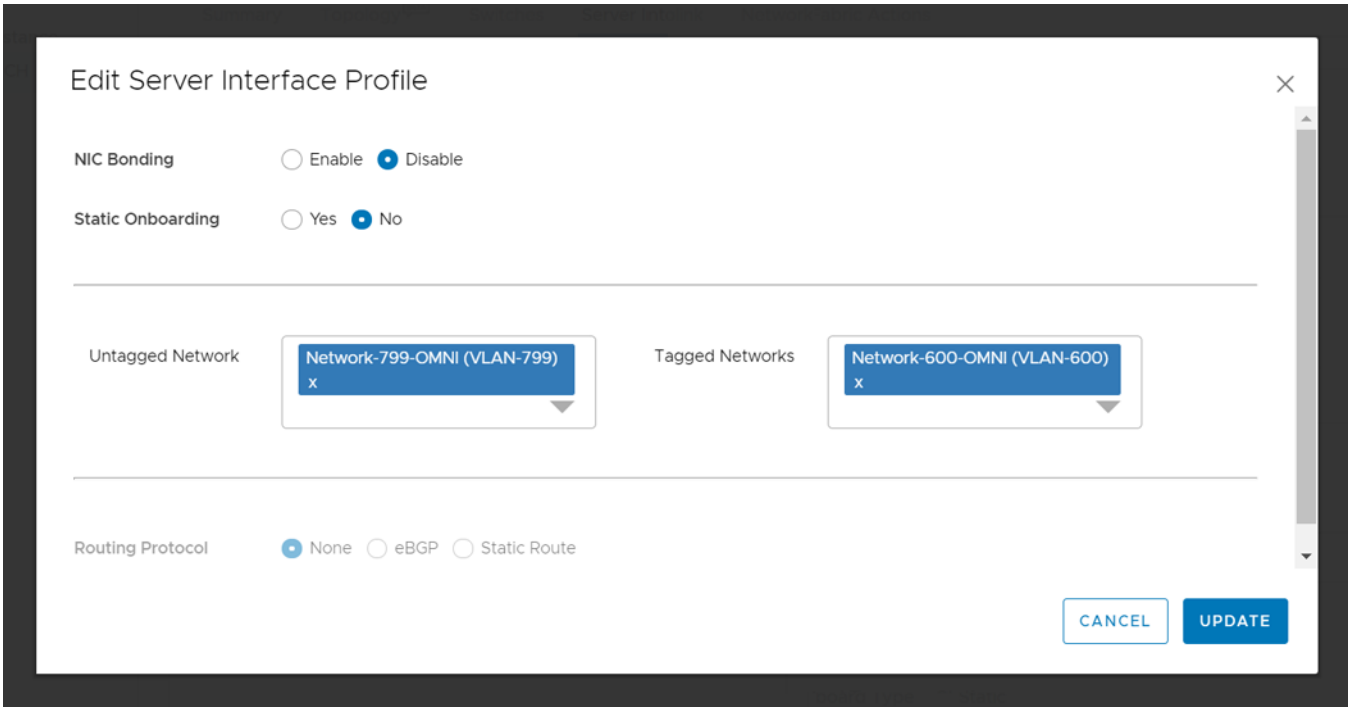

NOTE: You cannot select same network for both untagged and tagged networks.

(Optional) If the static onboarding is **Yes**, select **Leaf Node** and **Interface** (where the server interface is connected), select the **Routing Protocol**.

- (Optional) Select the **Routing Protocol** as **None**, and click **Update**.
- (Optional) Select the **Routing Protocol** as **eBGP**, enter the **ASN** and **IP address**, and click **Update**.
- (Optional) Select the **Routing Protocol** as **Static Route**, enter the **Network Address** and **Next-Hop Address**, and click **Update**.

NOTE: You cannot edit the server profile that is already configured in the system.

8. Click **Add Networks** to associate the networks that are part of the fabric for all the server interface profile. Select the networks for **Tagged Networks** and **Untagged Network** from the list, and click **Add**.

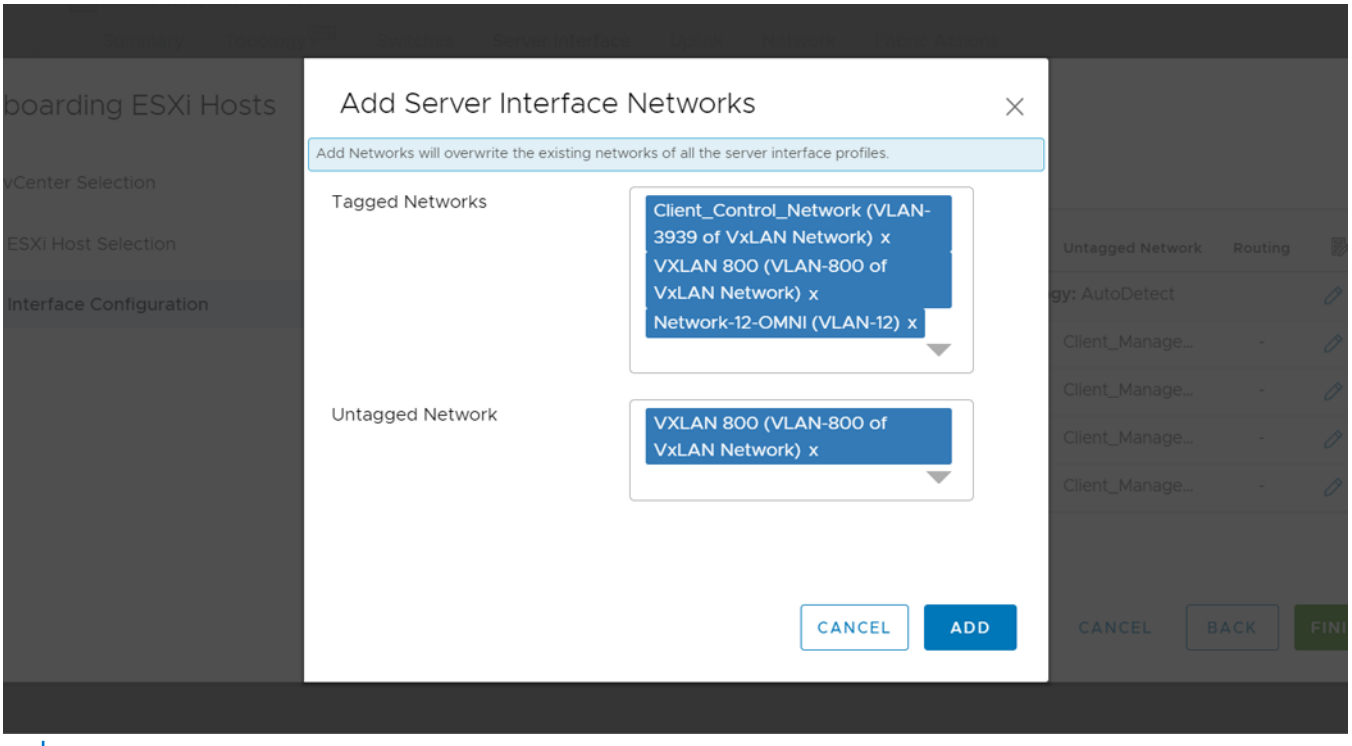

NOTE: Add networks overwrite the existing networks of all the server interface profiles.

9. Click **Finish** after all the configurations are complete.

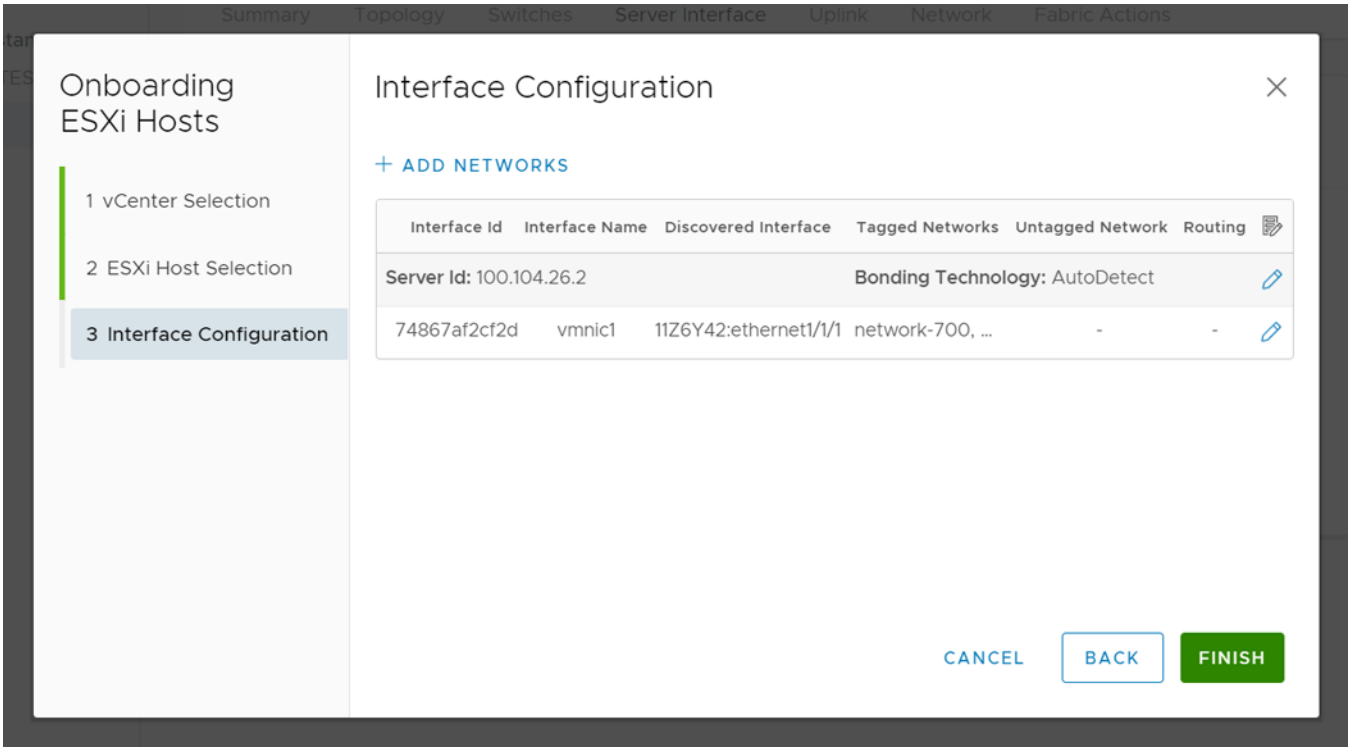

10. The system displays the server interface profile update success message.

Import SmartFabric discovered server interfaces

Automate onboarding of server interface profile by importing profiles that are discovered by SFS. Starting from 2.0 release, you can onboard unknown servers using OMNI.

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

- **Known server**—A known server is a host that sends a valid originator in Dell-specific (custom) TLVs in LLDP frame that is recognized by SFS. Following are the list of known servers that are discovered by SFS:
  - VxRail
  - PowerStore-X
  - PowerStore-T
- **Unknown server**—An unknown server is a host that does not send a valid originator in Dell-specific TLVs in LLDP frame.

1. From SmartFabric instance, select **Server Interface**.

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

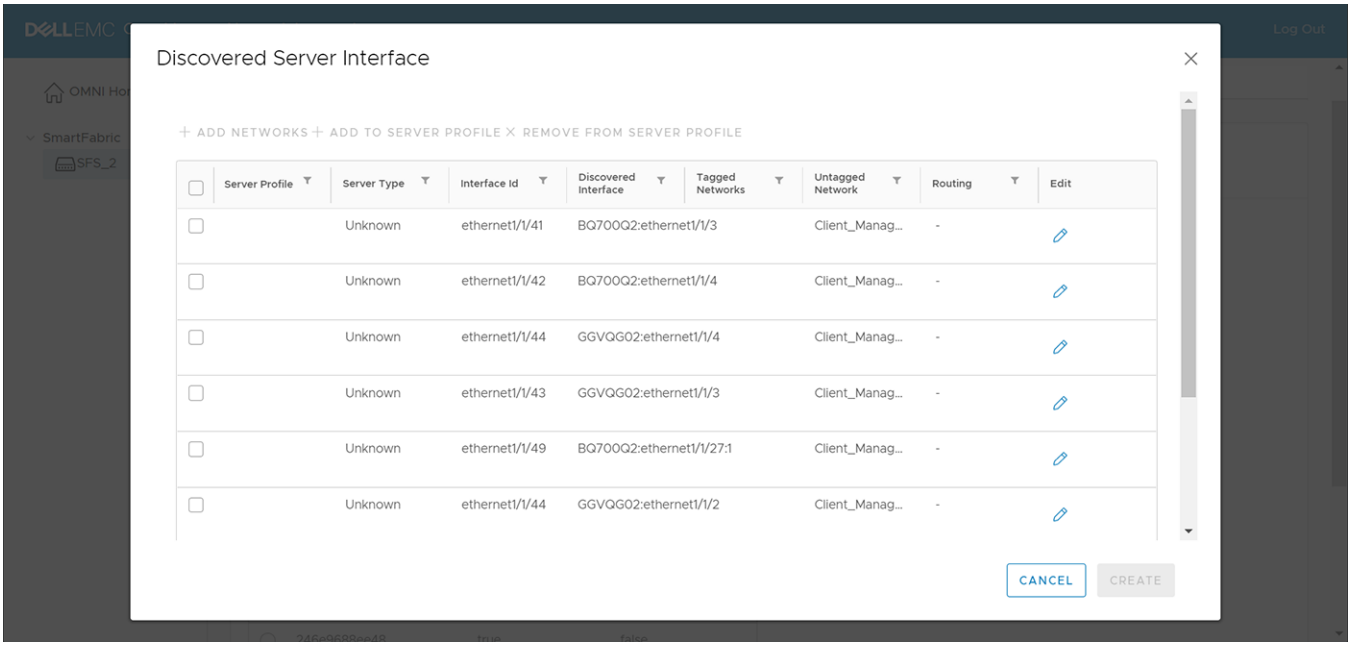

NOTE: The interface that is already associated with a server interface profile is not listed in the discovery table.

3. Edit the server profile information of each interface using the **Edit** option available at the end of each row.

   Edit the **NIC Bonding** configuration and **Static Onboarding**. If the static onboarding is **No**, select an **Untagged Network** and one or more **Tagged Networks** and click **Update**.

   NOTE: You cannot select same network for tagged and untagged network.

   (Optional) If static onboarding is **Yes**, select **Leaf Node** and **Interface** (where the server interface is connected), select the **Routing Protocol**.

   - (Optional) Select the **Routing Protocol** as **None**, and click **Update**.

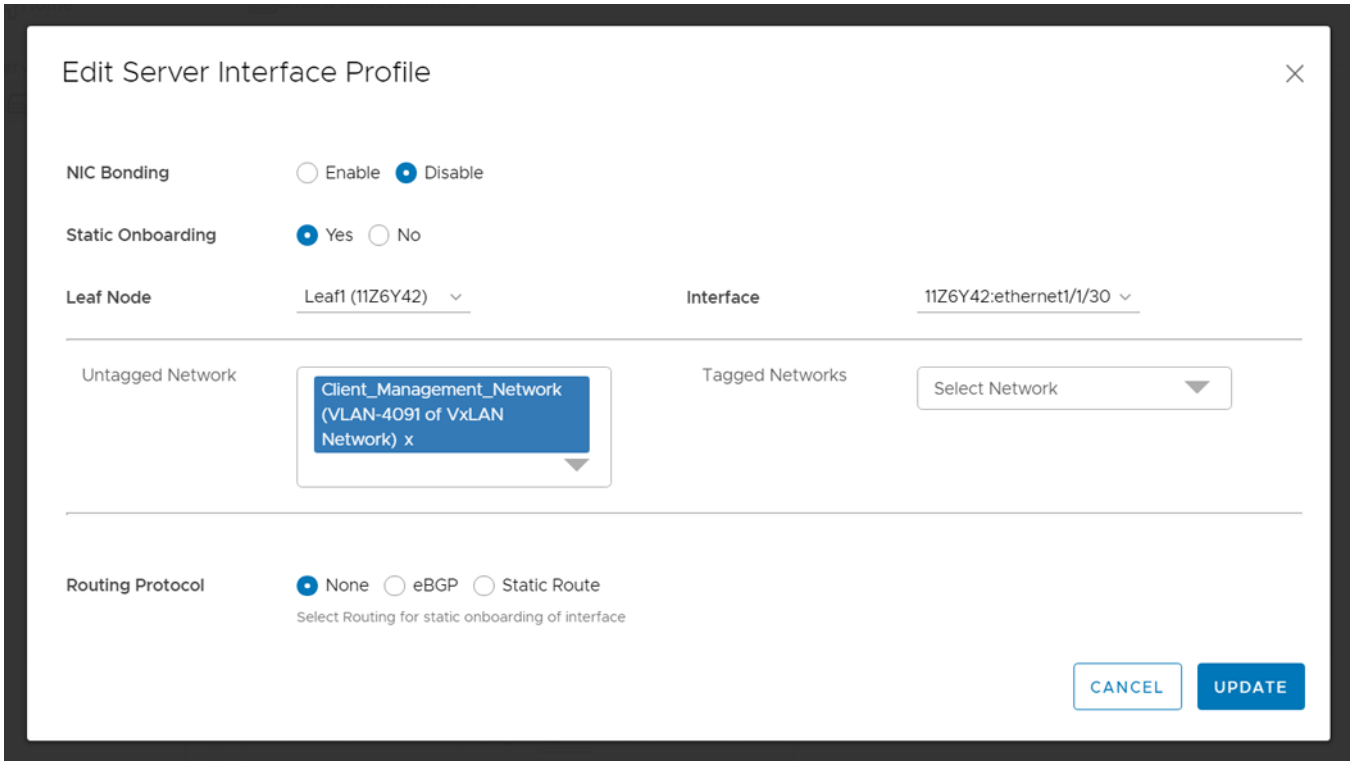

- (Optional) Select the **Routing Protocol** as **eBGP**, enter the **ASN** and **IP address**, and click **Update**.

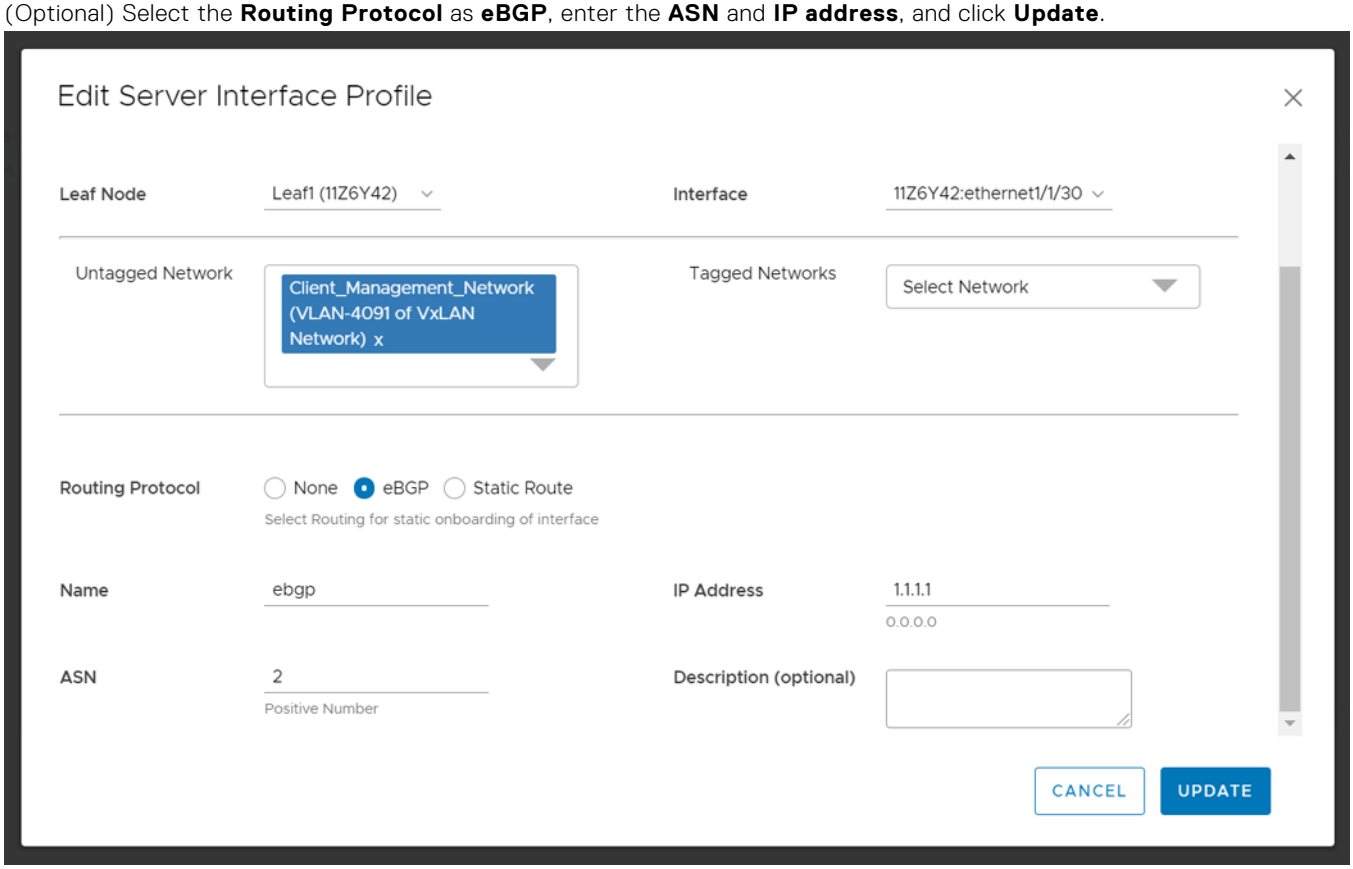

- (Optional) Select the **Routing Protocol** as **Static Route**, enter the **Network Address** and **Next-Hop Address**, and click **Update**.

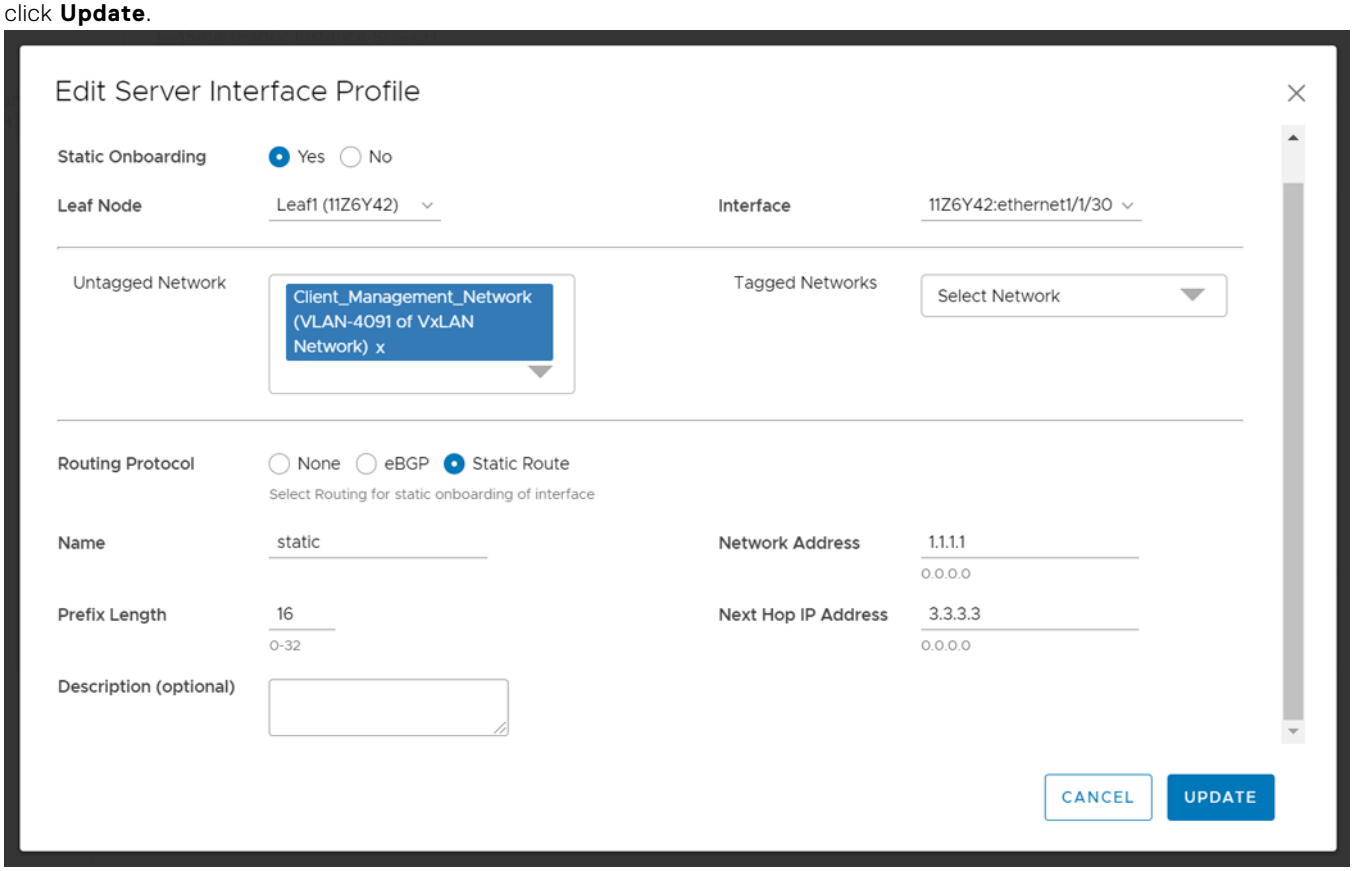

4. Select one or multiple discovered interfaces, add the service profile and networks, and click **Update**. For more information about adding server profile and networks, see *Add to Server Profile* and *Add networks* sections.

Add 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**.

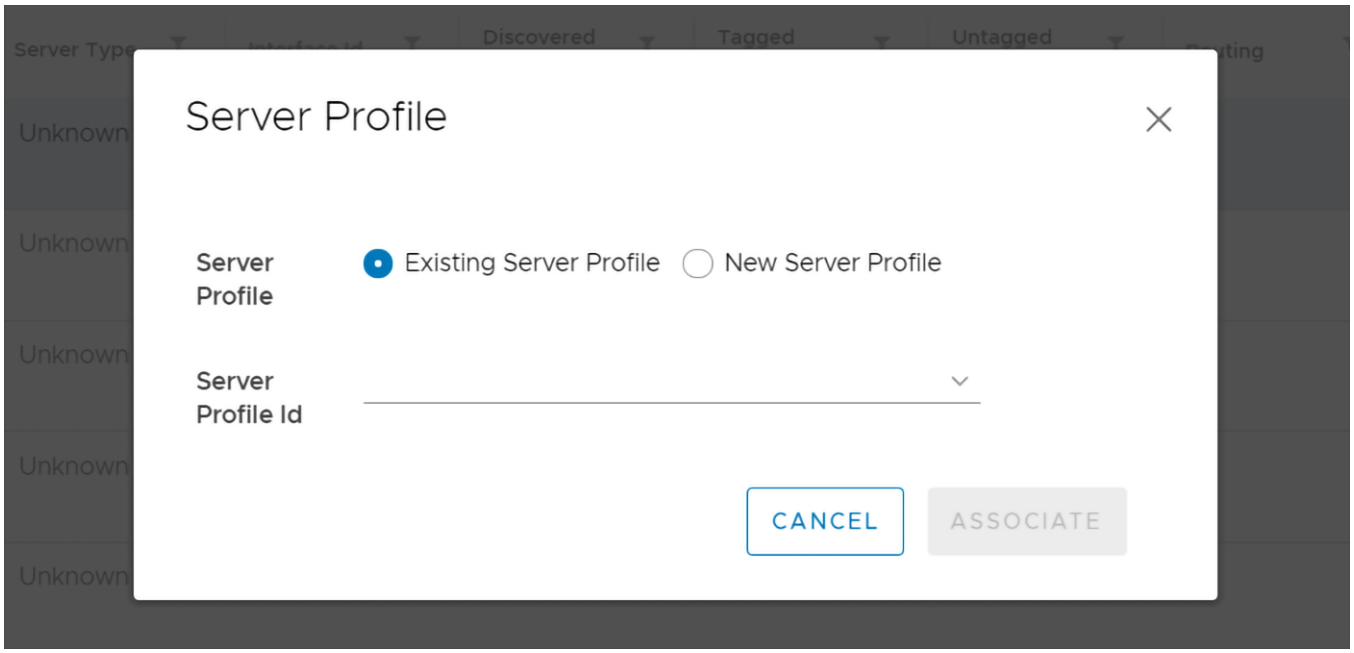

   - Select **New Server Profile**—Enter the **Server Profile Id** and **Bonding Type** to associate the interface with the new server profile, and click **Associate**.

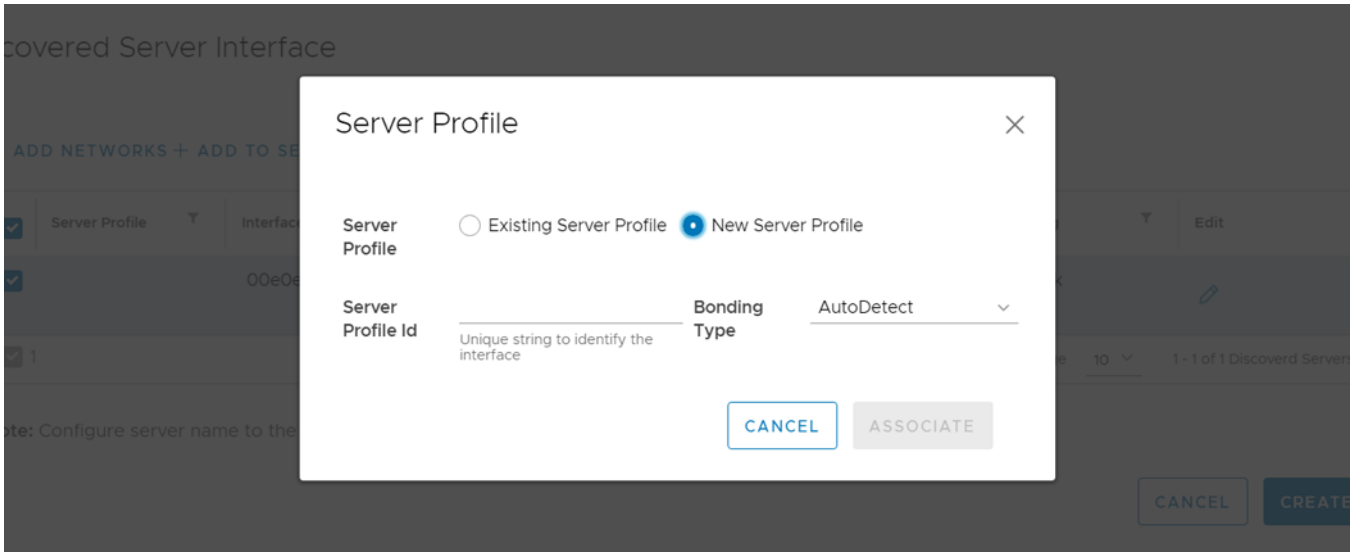

3. The system displays the server interface profile association success message.

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**.

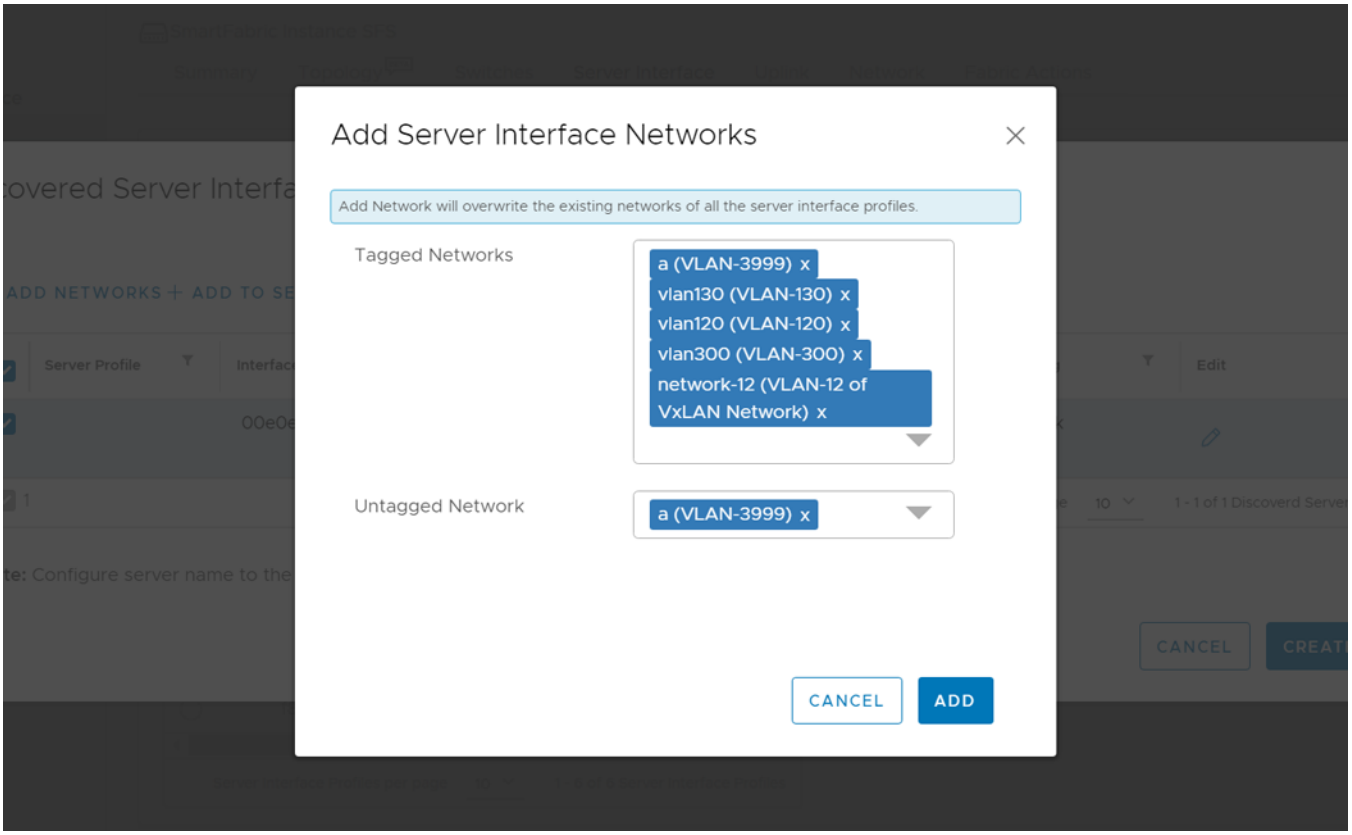

NOTE: Add networks overwrite the existing networks of all the server interface profiles.

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**.

Configure and manage uplinks

Configure an uplink and manage the uplinks that are available in the SmartFabric instance.

Using the **Uplinks** tab, you can:

- View the list of uplinks created in the SmartFabric instance.
- Create an uplink.
- Edit network and port configuration for an uplink.
- Delete a created uplink.

You can create uplinks with available interfaces which are not part of an existing uplink, server connected ports, part of a fabric automation, or jump port.

There are two types of uplinks—L2 and L3, and there are two types of L3 uplinks—L3 VLAN and L3 routed interface. Once you have created an uplink, you can then associate networks to the uplink and change or modify interfaces. These user-managed uplinks require configuration of networks through SmartFabric vCenter.

NOTE: If you delete an uplink, any unused networks and ports can be used for future use.

View Uplinks summary

From SmartFabric instance, select **Uplink**.

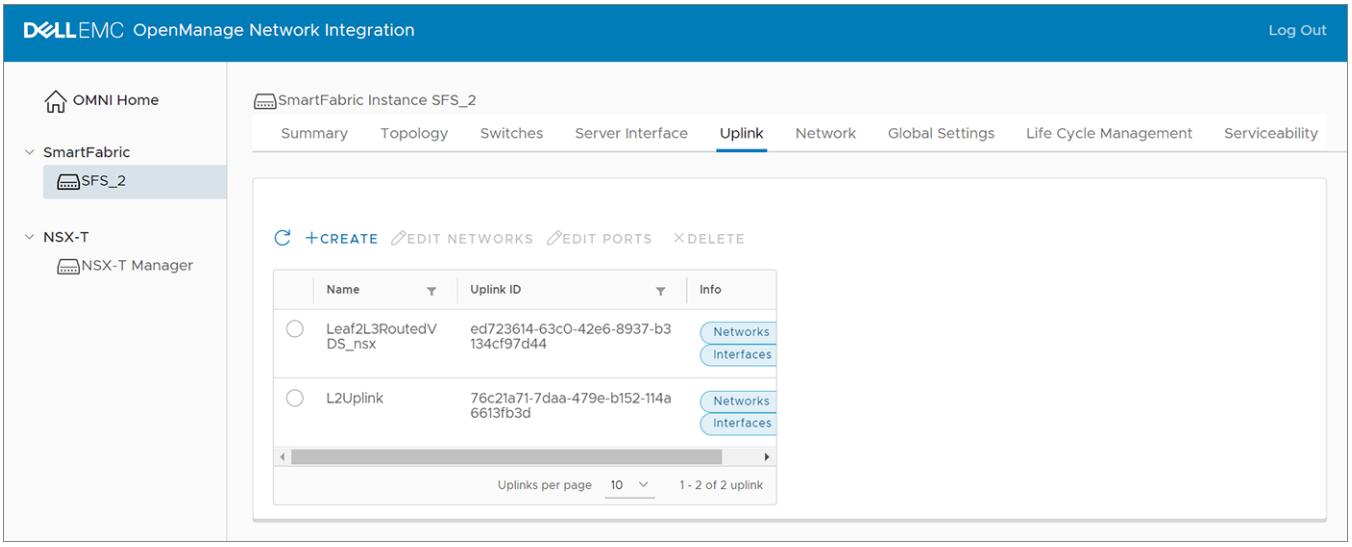

Create L2 Uplink

You can create an uplink by selecting the fabric with a unique name, and select the interfaces, and networks to create a user uplink.

1. Select the SmartFabric instance >**Uplink**, and click **Create**.

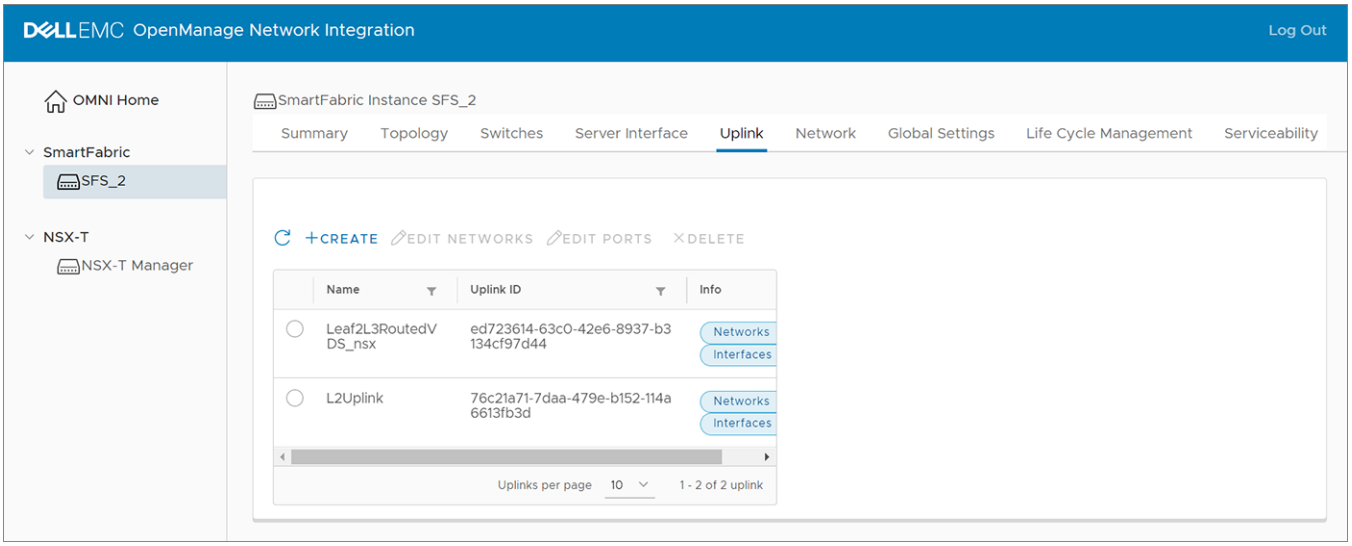

2. Enter the uplink port type as **L2**, a **Name**, an optional description, then click **Next**.

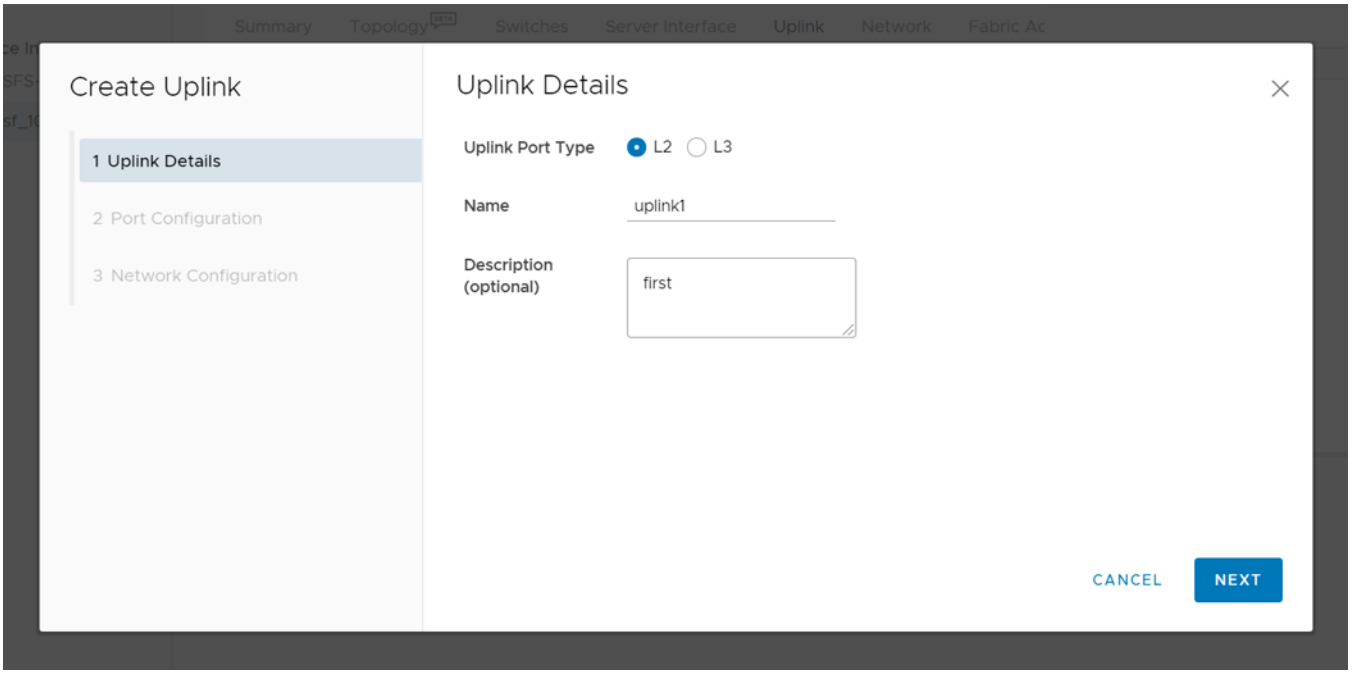

3. Enter the port configuration by selecting the rack to create the uplink on, select the interfaces, the **LAG Mode** (LACP or Static), then click **Next**.

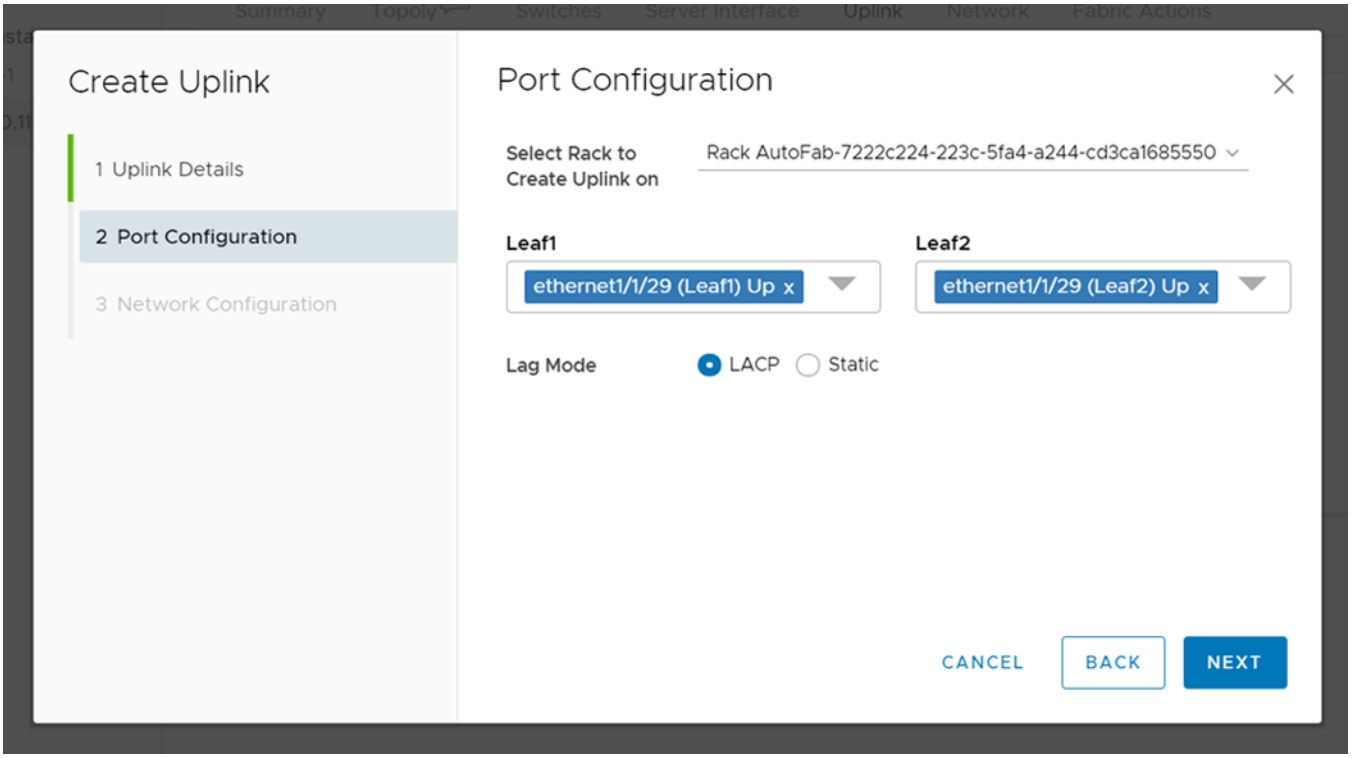

4. Select the untagged network, the OMNI network, and Select **Yes** or **No** to integrate the networks that are created automatically in the fabric through vCenter on this uplink.

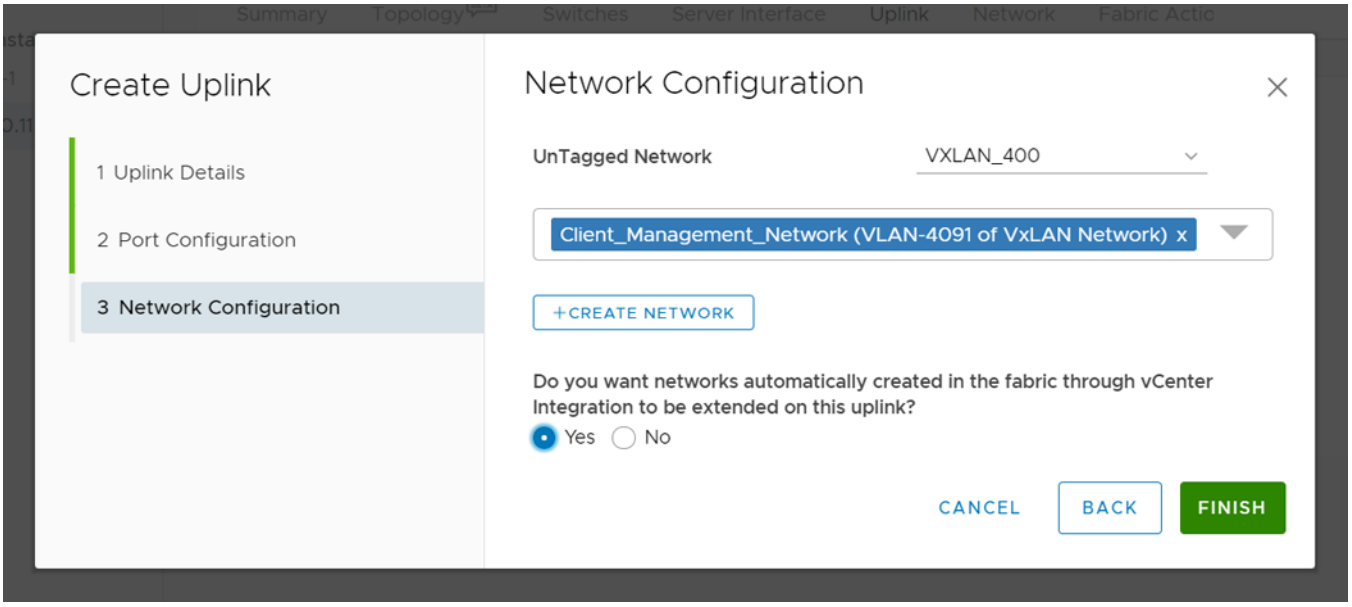

5. (Optional) Click **Create Network** to associate a network with the uplink. Enter the name of the network, optional description, and the VLAN number.

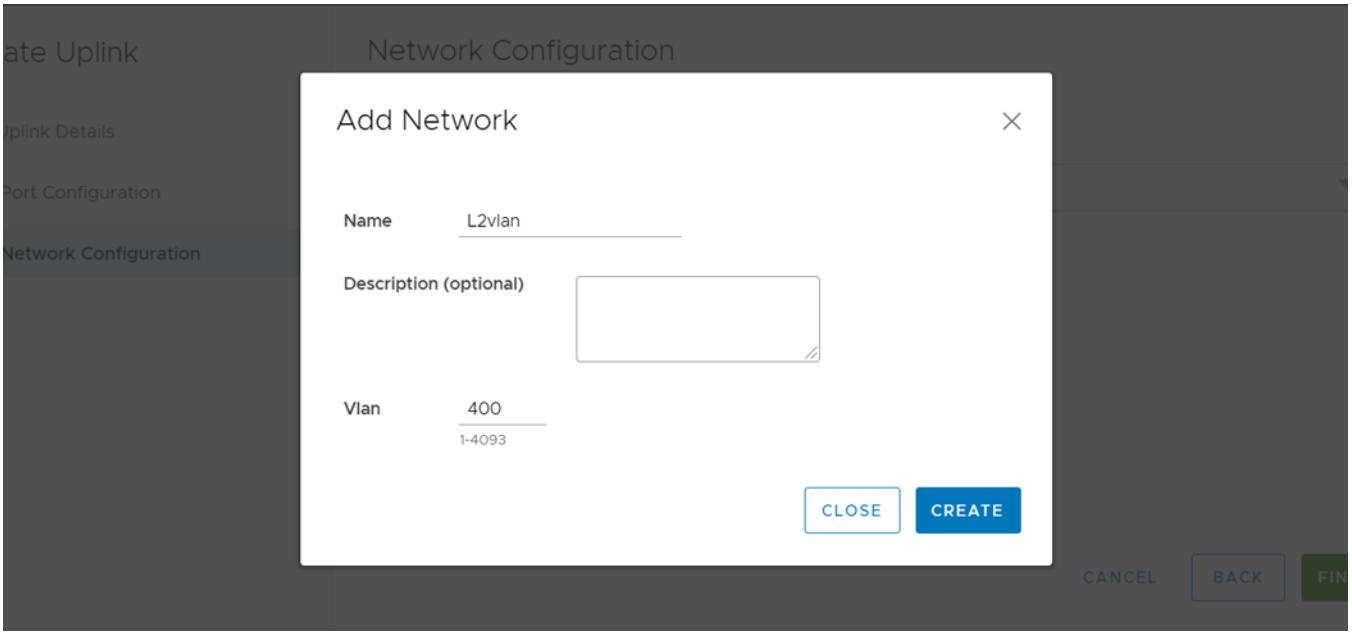

6. Click **Finish** to complete the L2 uplink creation.
7. The system displays user uplink creation success message.

Create L3 Uplink

Create an L3 uplink of L3 VLAN or L3 routed interface types.

Create L3 VLAN uplink

1. Select the SmartFabric instance > **Uplink**, and click **Create**.

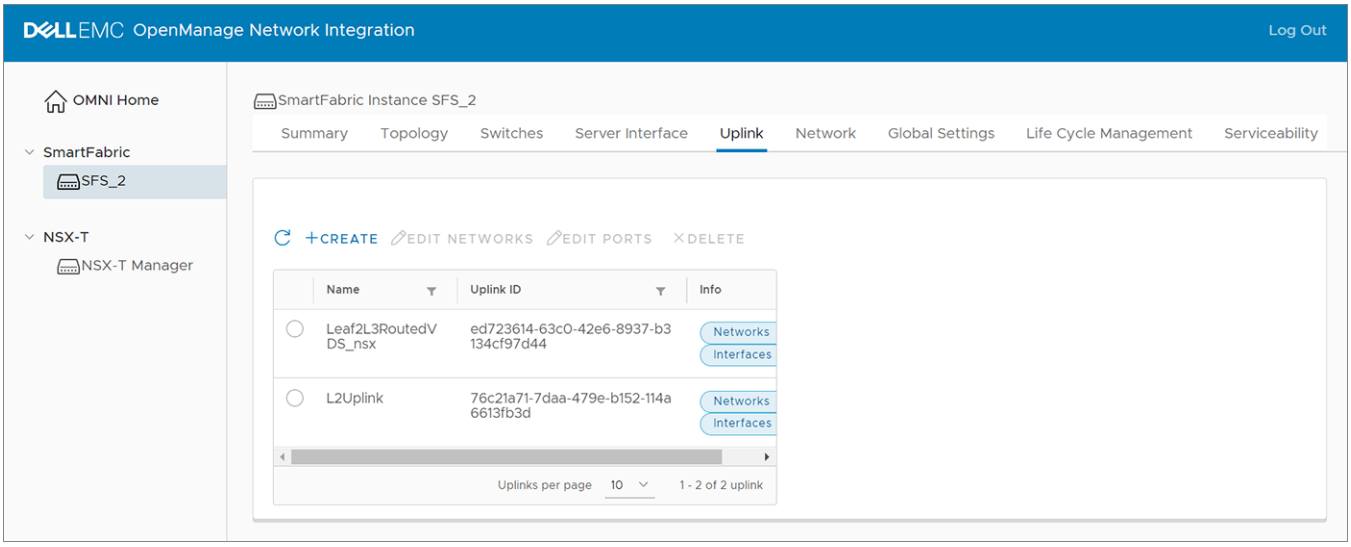

2. Select **L3** for the uplink port type, select **L3 VLAN**, enter the **name** for the uplink, and optional description, then click **Next**.

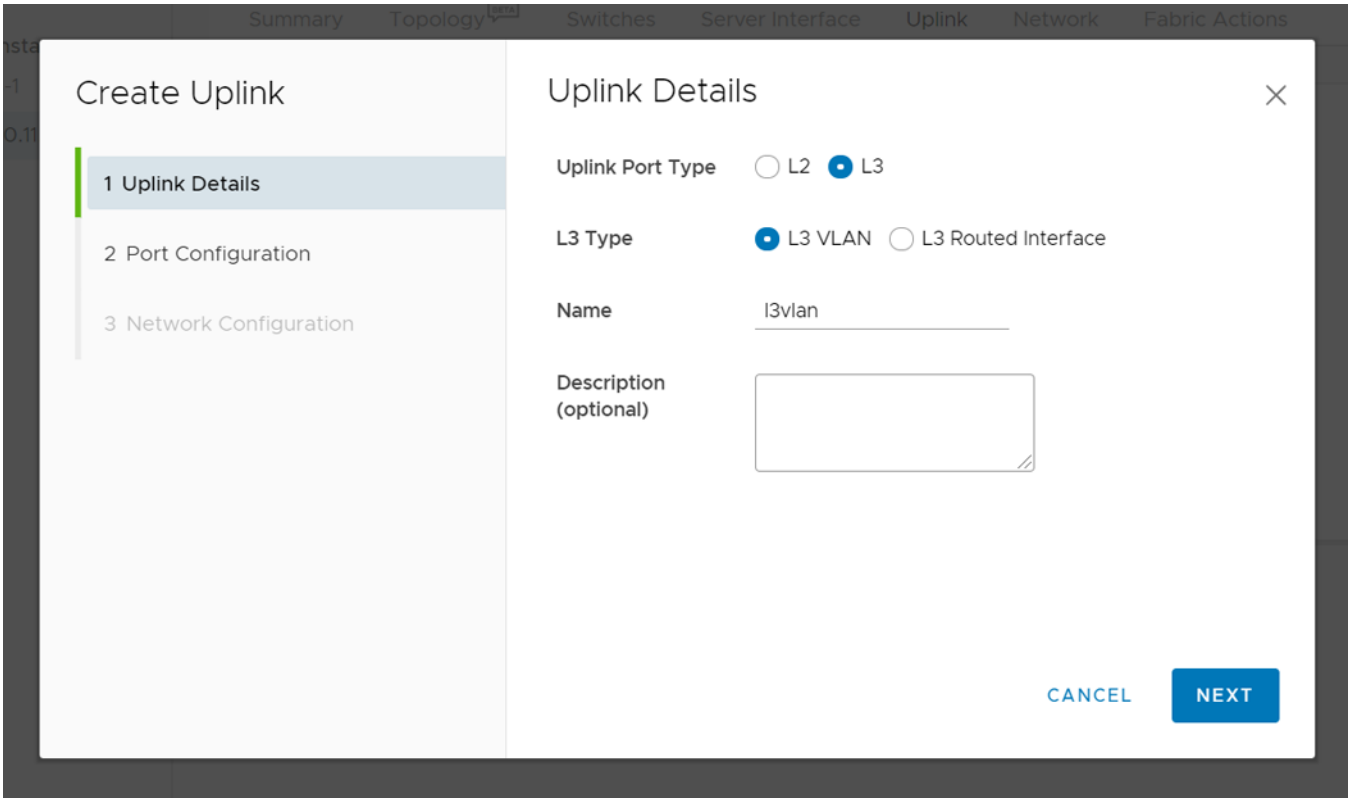

3. Select the **Switch group** (Leaf or Spine), select the **rack** to create the uplink on, select the **interfaces**, select **LACP** for the LAG mode, then click **Next**.

   Leaf:

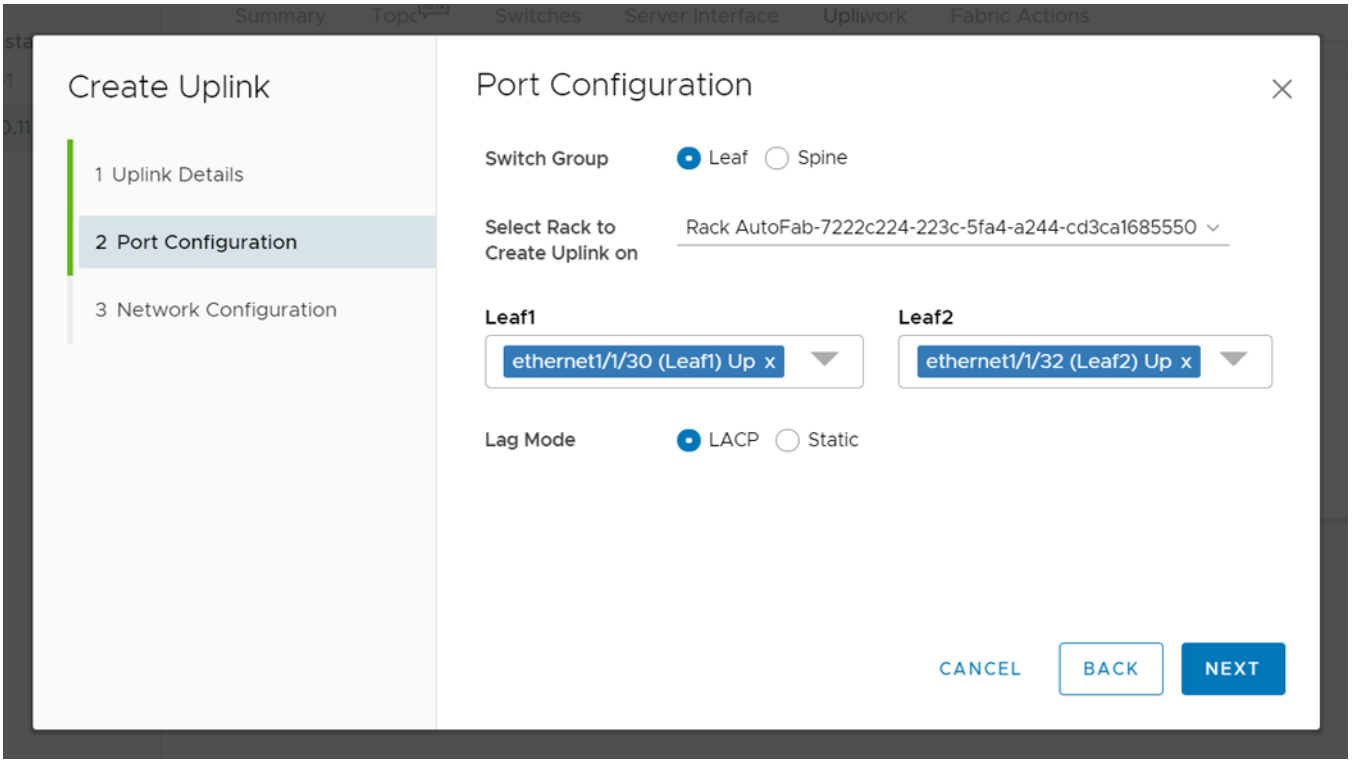
Summary
Topo
Switches
Server Interface
Uplwork
Fabric Actions
Create Uplink
1 Uplink Details
2 Port Configuration
3 Network Configuration
Port Configuration
Switch Group
Leaf
Spine
Select Rack to Create Uplink on
Rack AutoFab-7222c224-223c-5fa4-a244-cd3ca1685550
Leaf1
ethernet1/1/30 (Leaf1) Up x
Leaf2
ethernet1/1/32 (Leaf2) Up x
Lag Mode
LACP
Static
CANCEL
BACK
NEXT

Spine:

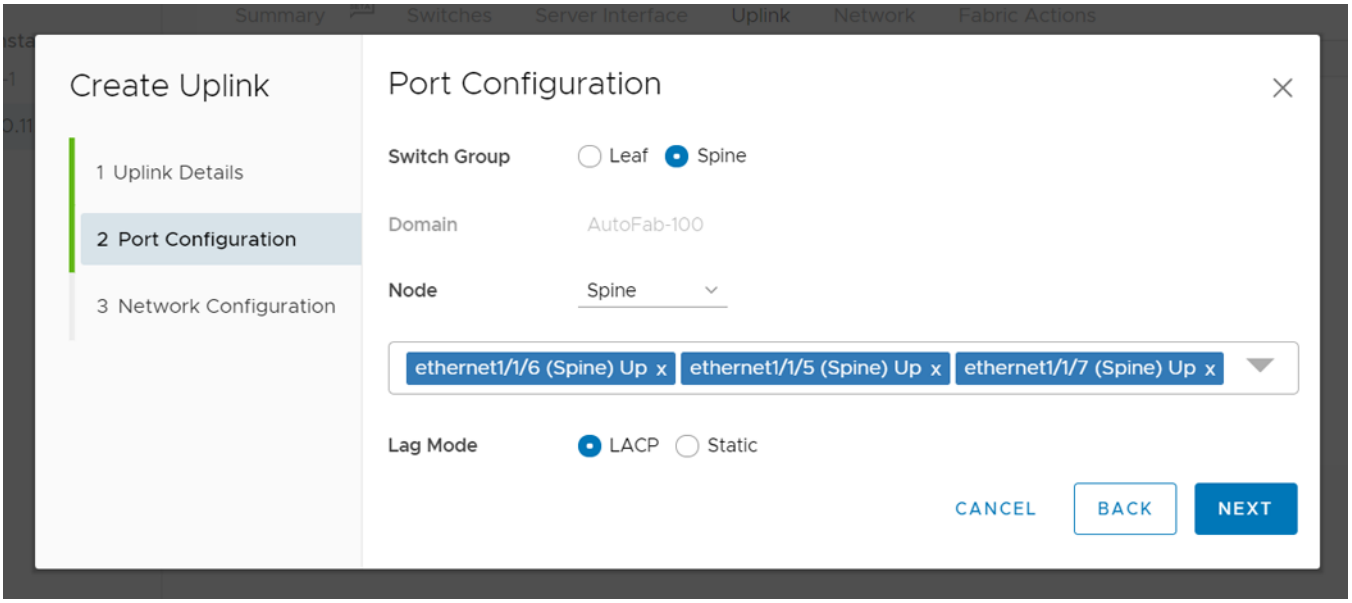
Summary
Switches
Server Interface
Uplink
Network
Fabric Actions
Create Uplink
1 Uplink Details
2 Port Configuration
3 Network Configuration
Port Configuration
Switch Group
Leaf
Spine
Domain
AutoFab-100
Node
Spine
ethernet1/1/6 (Spine) Up x
ethernet1/1/5 (Spine) Up x
ethernet1/1/7 (Spine) Up x
Lag Mode
LACP
Static
CANCEL
BACK
NEXT

4. Select **UnTagged** network, select the **OMNI network**, enter an optional description, select either **eBGP** or **Static Route** for the routing protocol, enter the routing policy information, then click **Finish**.

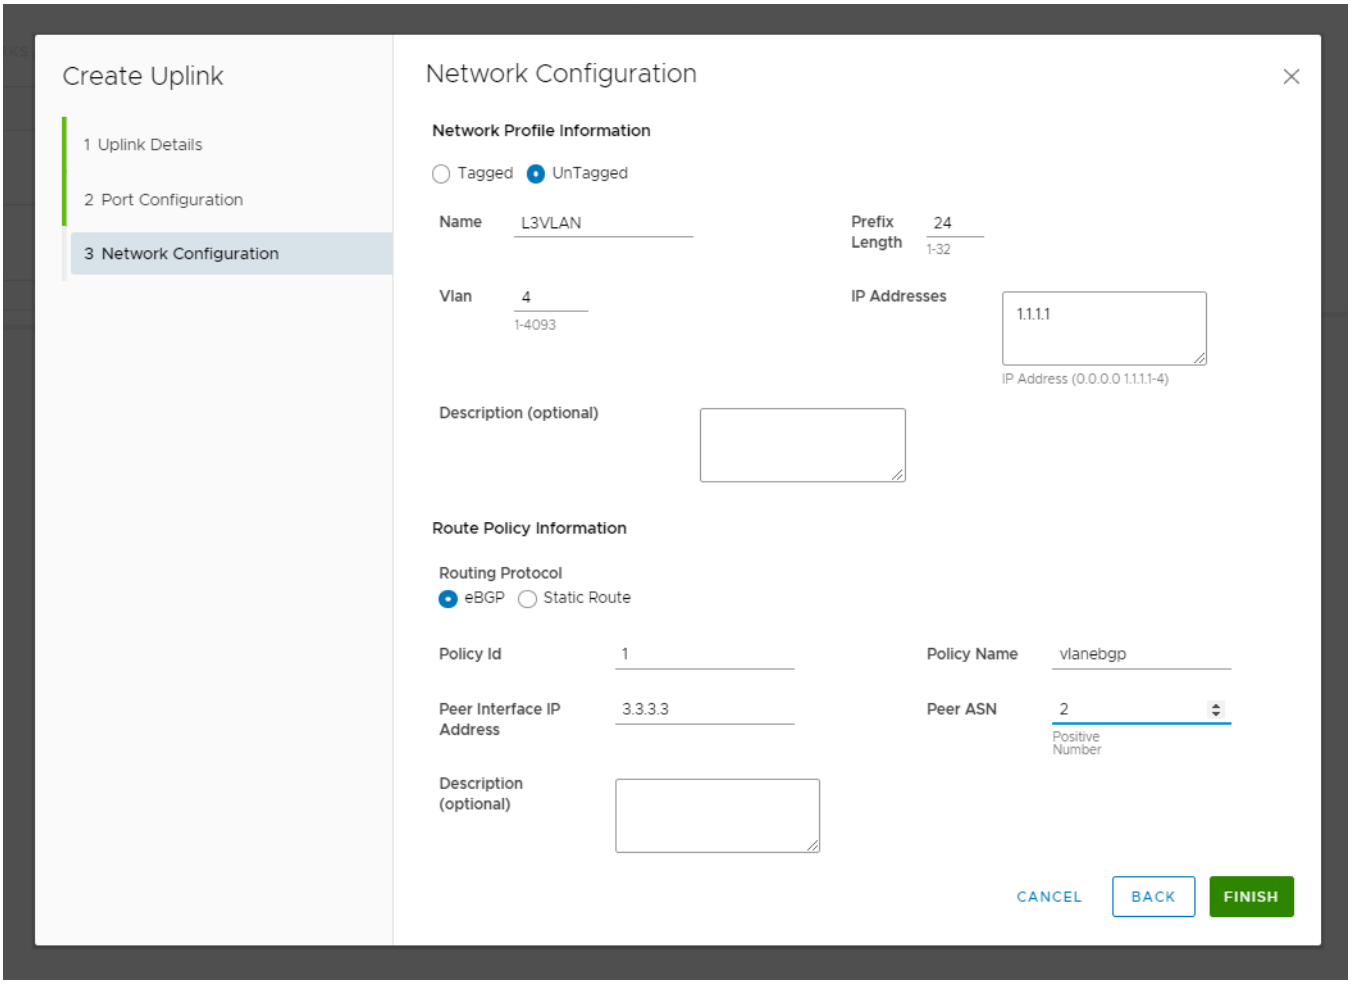

5. A route is associated with the nodes that are configured in the port configuration. The system displays uplink creation success message.

Create L3 routed interface uplink

1. Select the SmartFabric Instance > **Uplink**, and click **Create**.
2. Select **L3 routed interface**, enter the **Uplink name**, and optional description, then click **Next**.

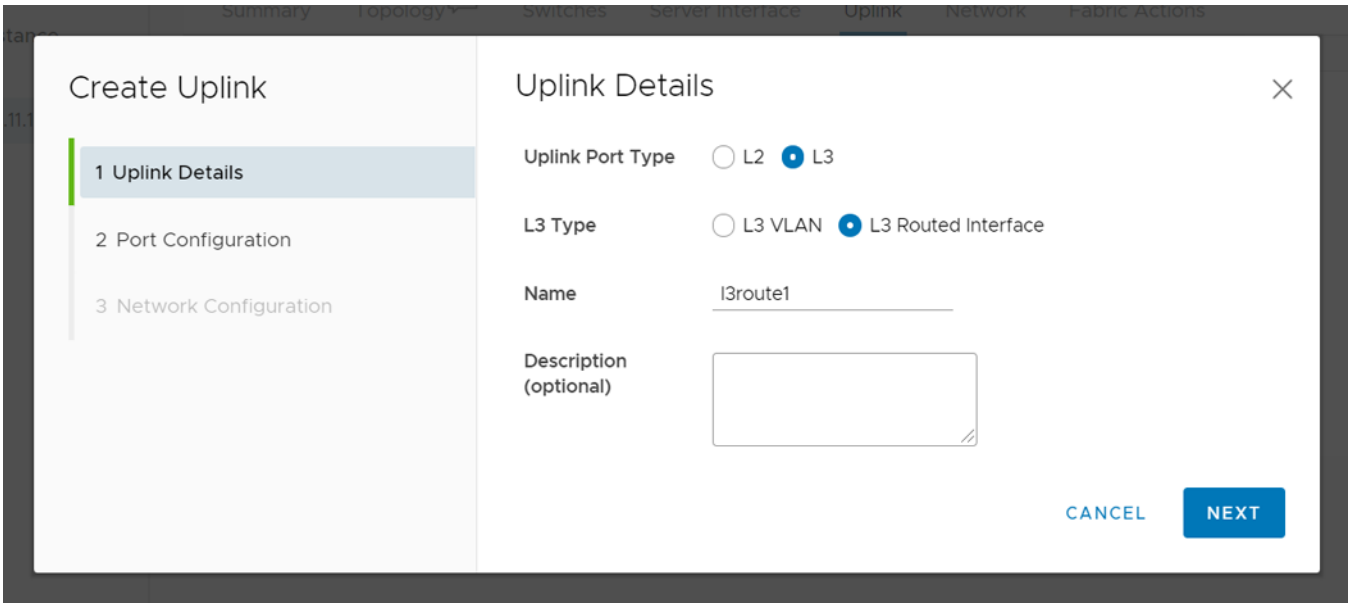

3. Select the Switch group (Leaf or Spine), the **rack** to create the uplink on, select the **interfaces**, then click **Next**.

Leaf:

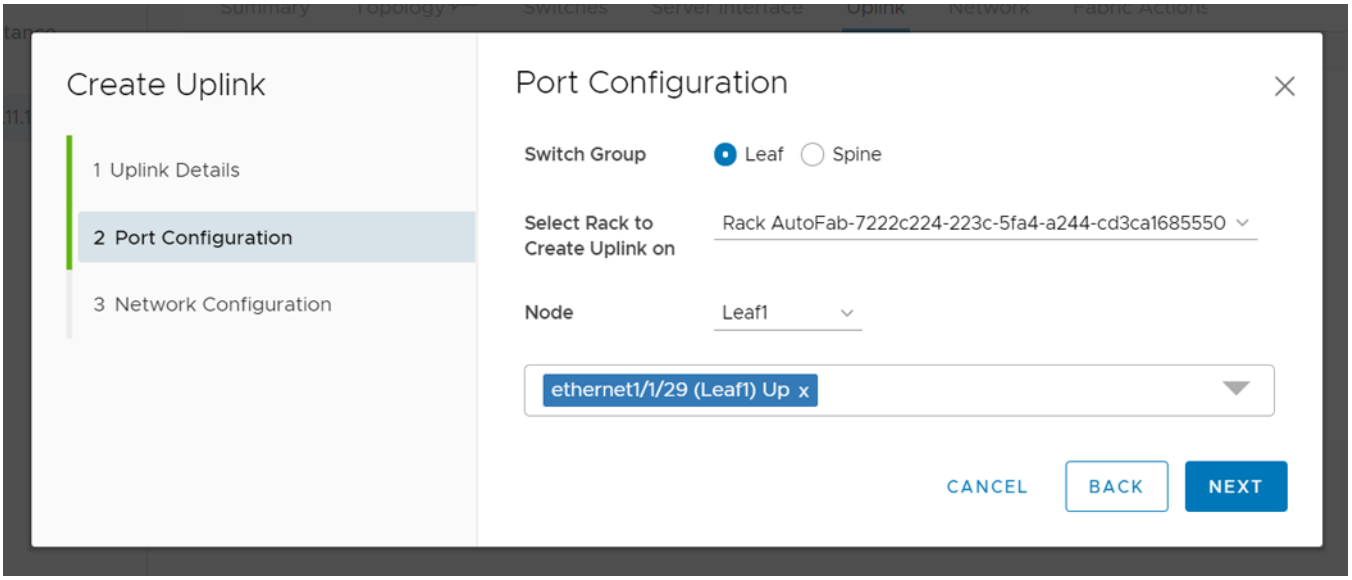
Create Uplink
1 Uplink Details
2 Port Configuration
3 Network Configuration
Port Configuration
Switch Group
Leaf
Spine
Select Rack to Create Uplink on
Rack AutoFab-7222c224-223c-5fa4-a244-cd3ca1685550
Node
Leaf1
ethernet1/1/29 (Leaf1) Up x
CANCEL
BACK
NEXT

Spine:

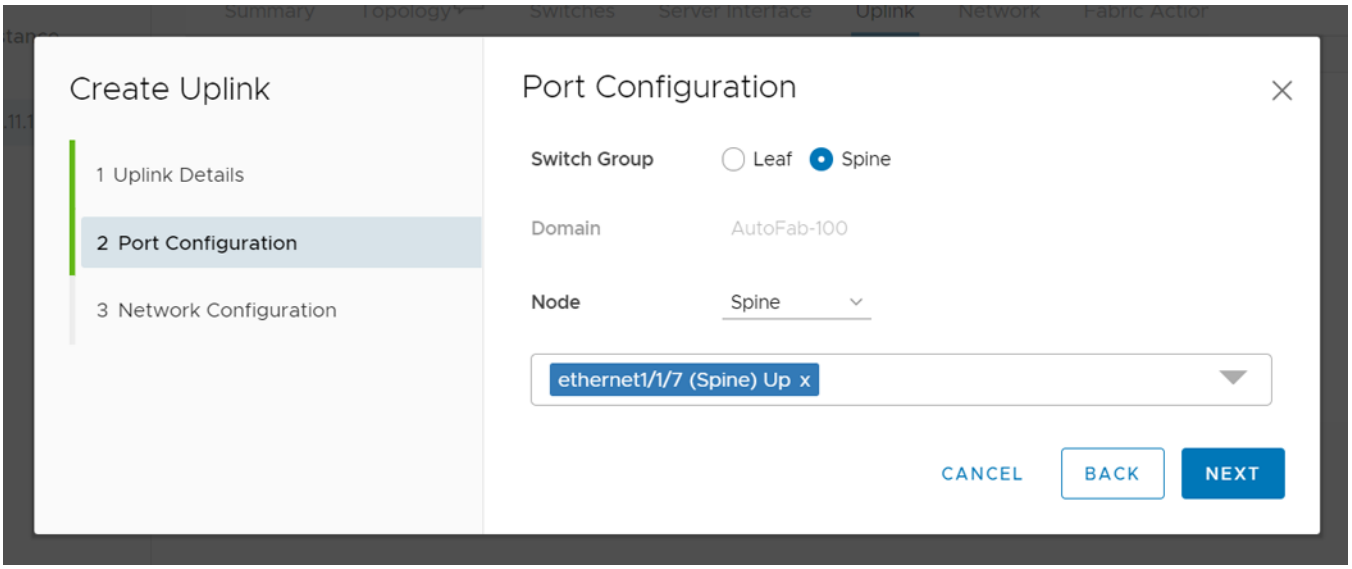
Create Uplink
1 Uplink Details
2 Port Configuration
3 Network Configuration
Port Configuration
Switch Group
Leaf
Spine
Domain
AutoFab-100
Node
Spine
ethernet1/1/7 (Spine) Up x
CANCEL
BACK
NEXT

4. Enter the network profile information and routing policy information for the uplinks, then click **Finish**.

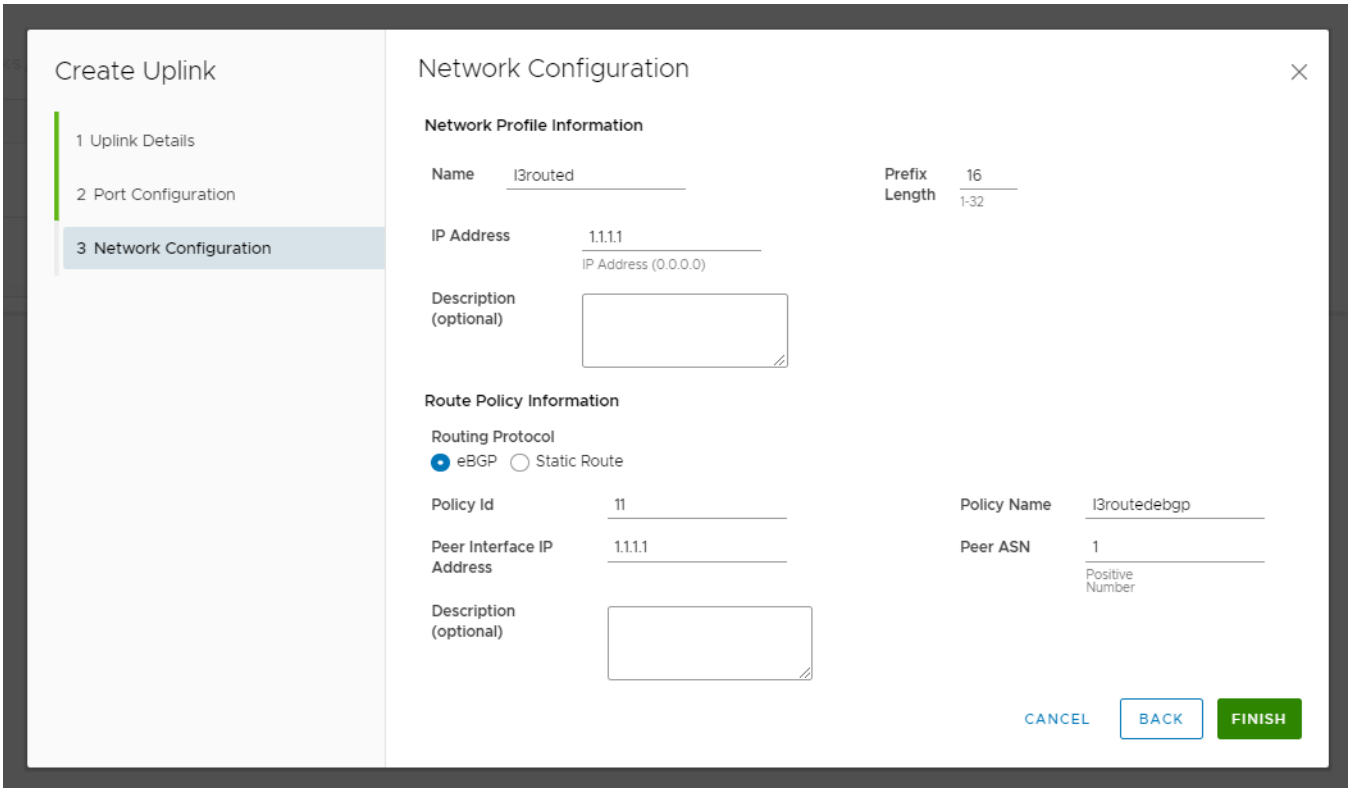

5. The system displays L3 routed uplink creation success message.

Edit networks and ports in an uplink

You can edit the network and port configuration for an uplink, and also view the detailed information of the uplink.

Select the uplink from the displayed list to view the details of the uplink on the right.

Edit networks

1. Select the uplink from the list, and click **Edit Networks**.

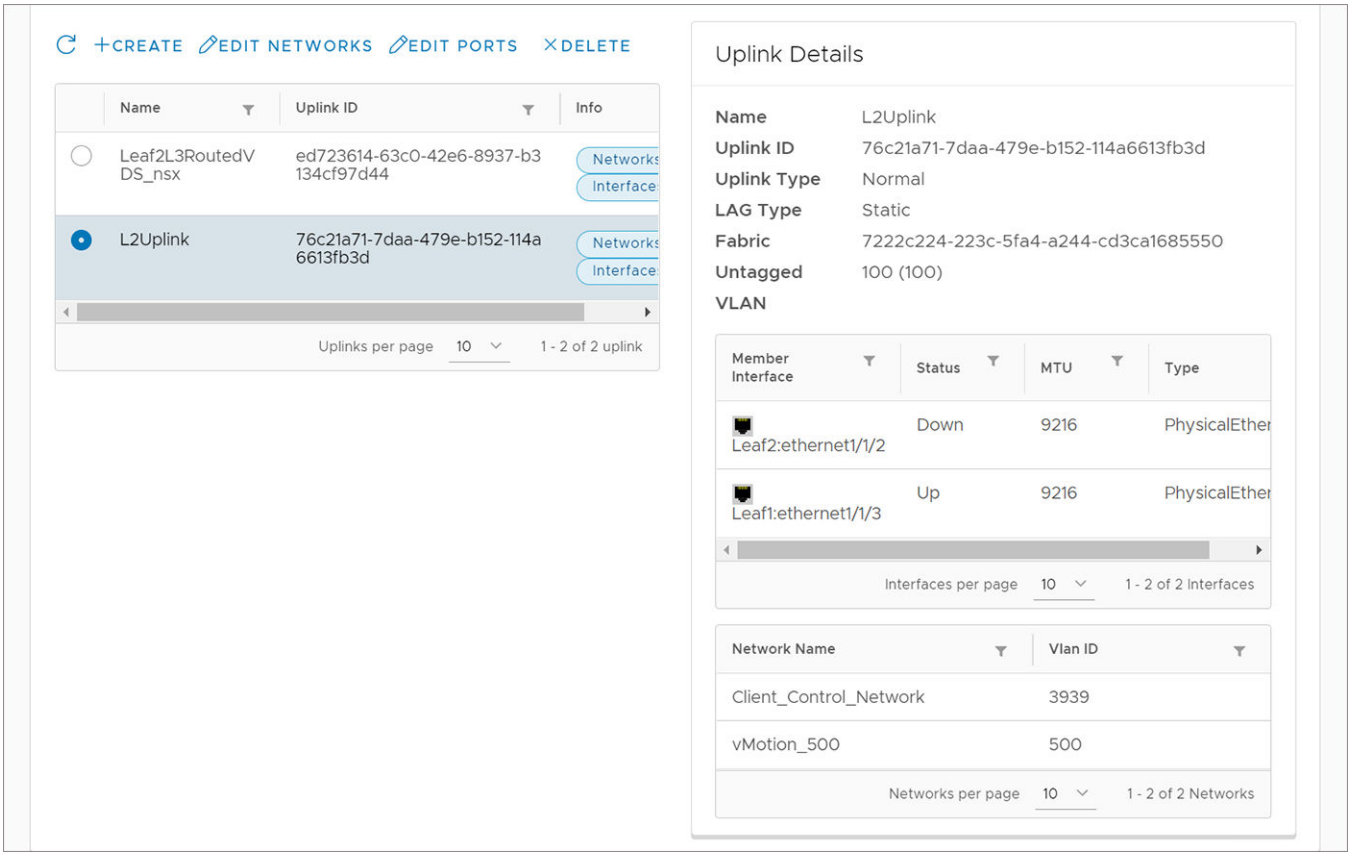

2. Edit the **Untagged Network** associated with the uplink, and click **Update**.

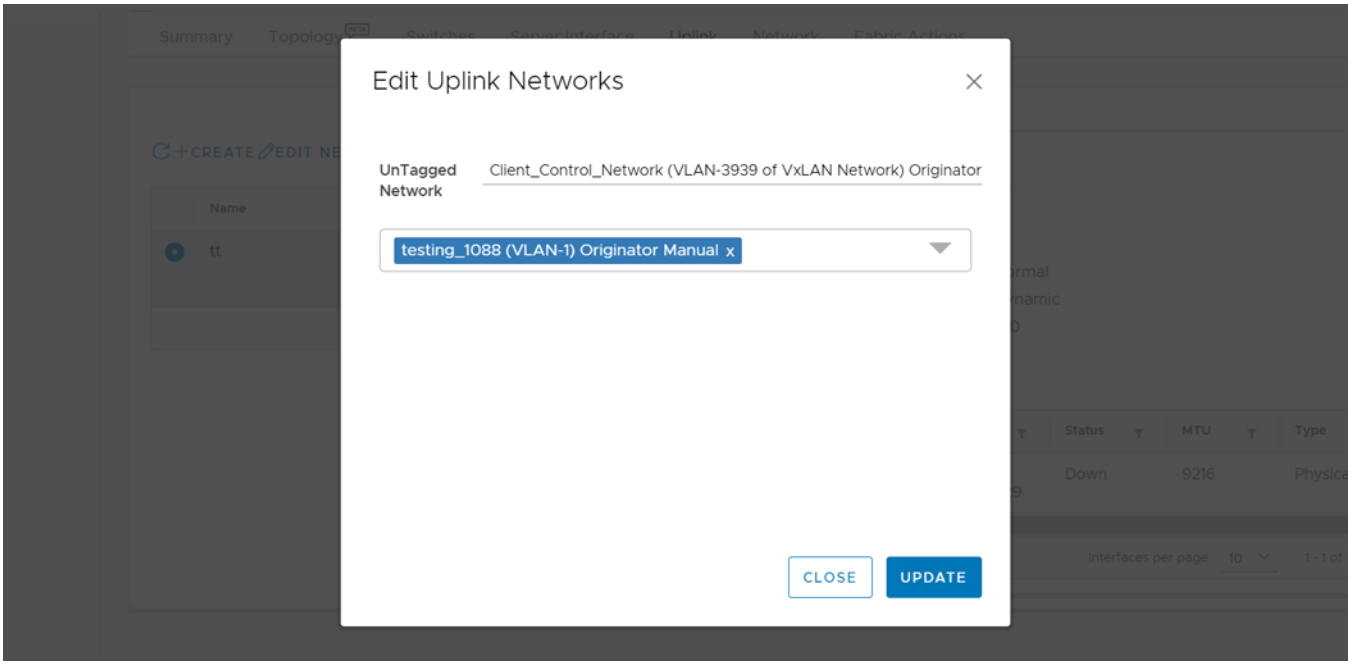

3. The system displays the uplink interface edit success message.

Edit ports

1. Select the fabric uplink from the list, and click **Edit Ports**.

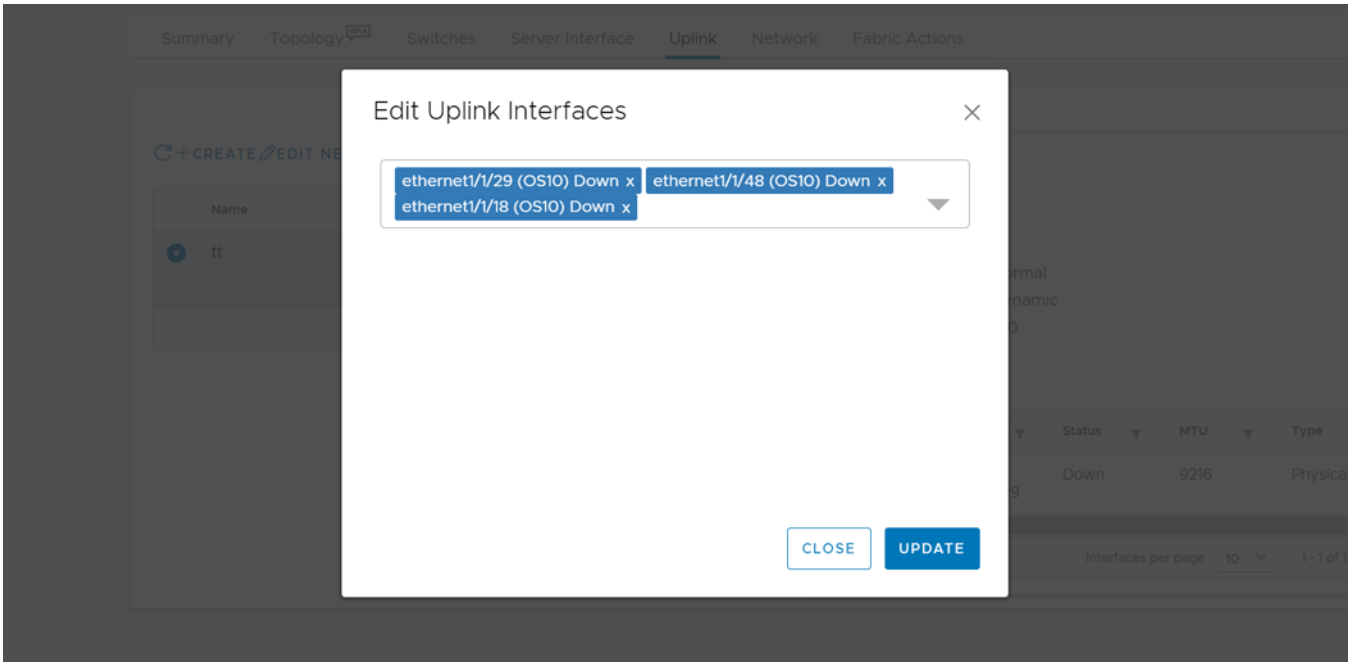

2. Edit the networks associated with uplink interfaces, and click **Update**.
3. The system displays the uplink interface edit success message.

Delete an uplink

You can delete a user-created uplink. To delete:

1. Select the uplink from the displayed list, and click **Delete**.

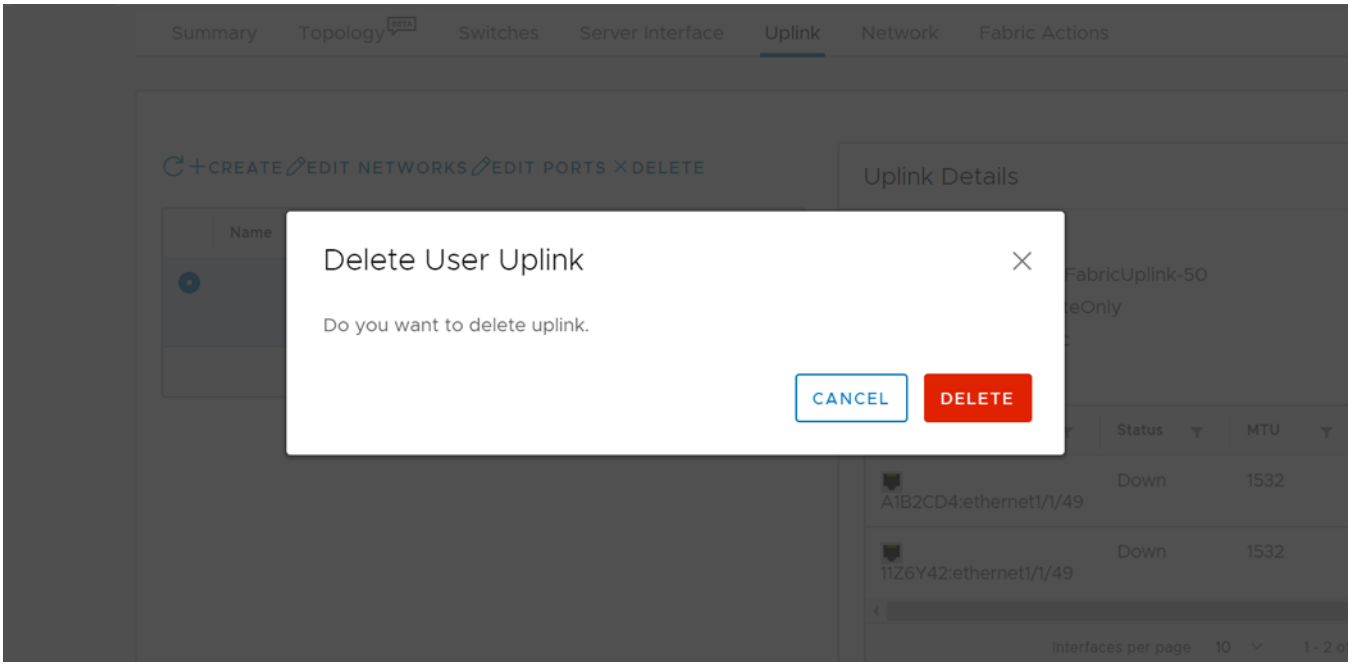

2. Click **Delete** to confirm.

Configure networks and routing configuration

You can set up networks and routing configuration.

NOTE: Networks that are created by the OMNI user interface are considered *Manual*.

The OMNI vCenter `PortGroup` VLAN automation process does not add *Manual* networks to auto uplinks, and does not remove them from SmartFabric. Add *Manual* networks to uplinks using the OMNI portal if needed. The OMNI VLAN automation process uses *Manual* networks for `ServerInterfaces`. If you are using the VLANs for the OMNI registered vCenter `PortGroup`, it is not recommended to use the OMNI portal to create a network. OMNI automation manages those VLANs or networks by itself. For complete information, see OMNI vCenter integration.

You can configure five types of networks including multi rack L3 VLAN, general purpose, VXLAN networks (for L2 and L3 profiles), VLAN networks (for L2 and L3 profiles), and L3 routed interfaces (for L3 profiles only).

Configure networks

You can manage general purpose, VXLAN and VLAN networks, and L3 routed interfaces.

From SmartFabric, select the instance > **Network**. From **Network** tab, you can create, edit, and delete general purpose, VXLAN and VLAN networks, and L3 routed interfaces.

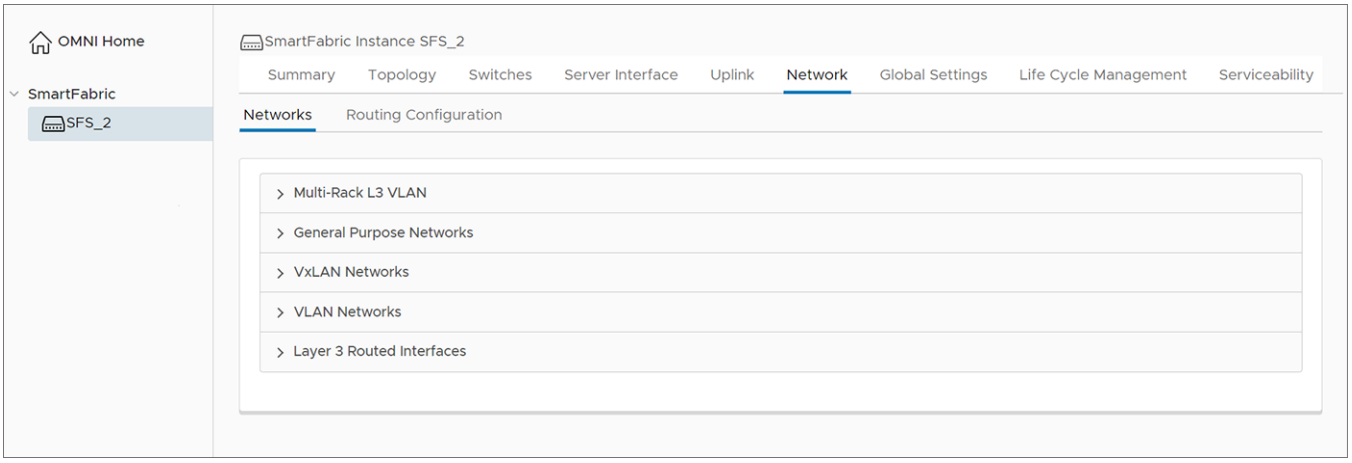

Configure multi rack L3 VLAN

Starting from 2.0 release, OMNI allows you to 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 multi rack L3 VLAN networks from OMNI.

This feature is used as part NSX-T workflow. As part of automation, OMNI creates all the NSX-T networks as multi rack L3 VLAN networks. You can edit and provide the Layer 3 details to the NSX-T networks. For more information, see OMNI support for NSX-T.

Create multi rack L3 VLAN

1. From SmartFabric instance, select **Networks** > **Multi-Rack L3 Networks**, and click **Create**.

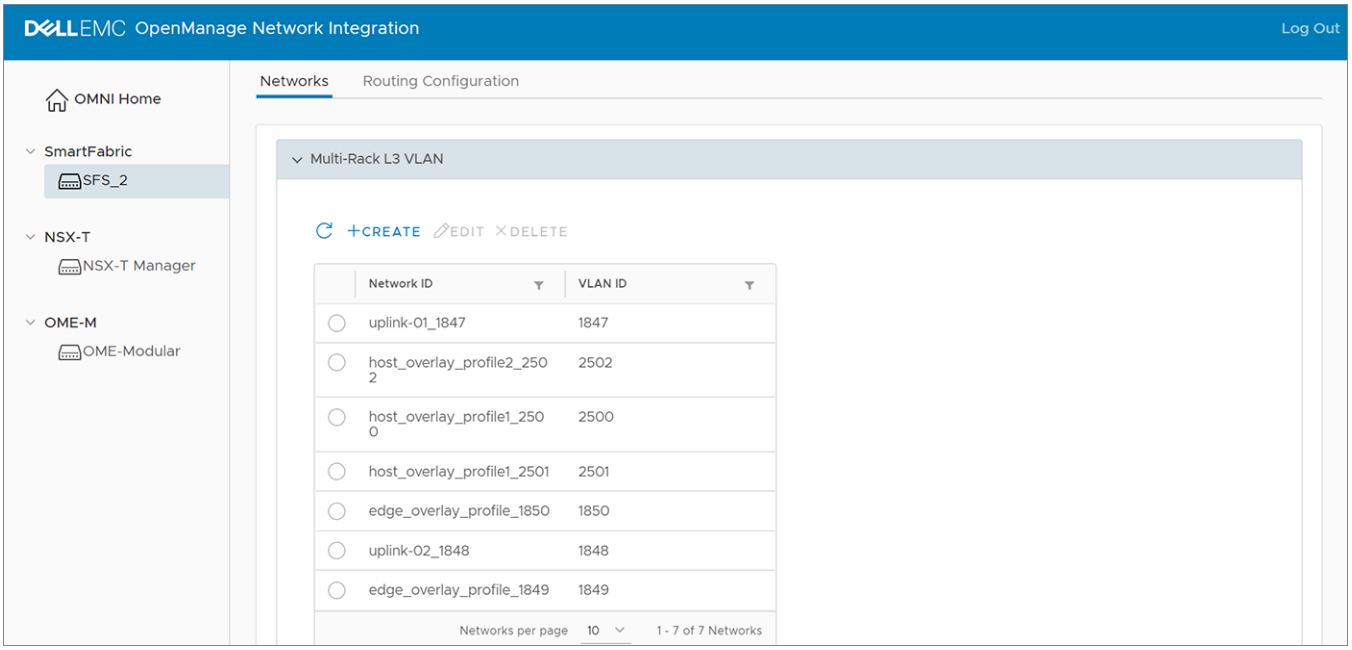

2. Enter the network ID, name, VLAN number, IP addresses, description, gateway IP address, and helper addresses for each rack available in the SmartFabric cluster. Click **Create**.

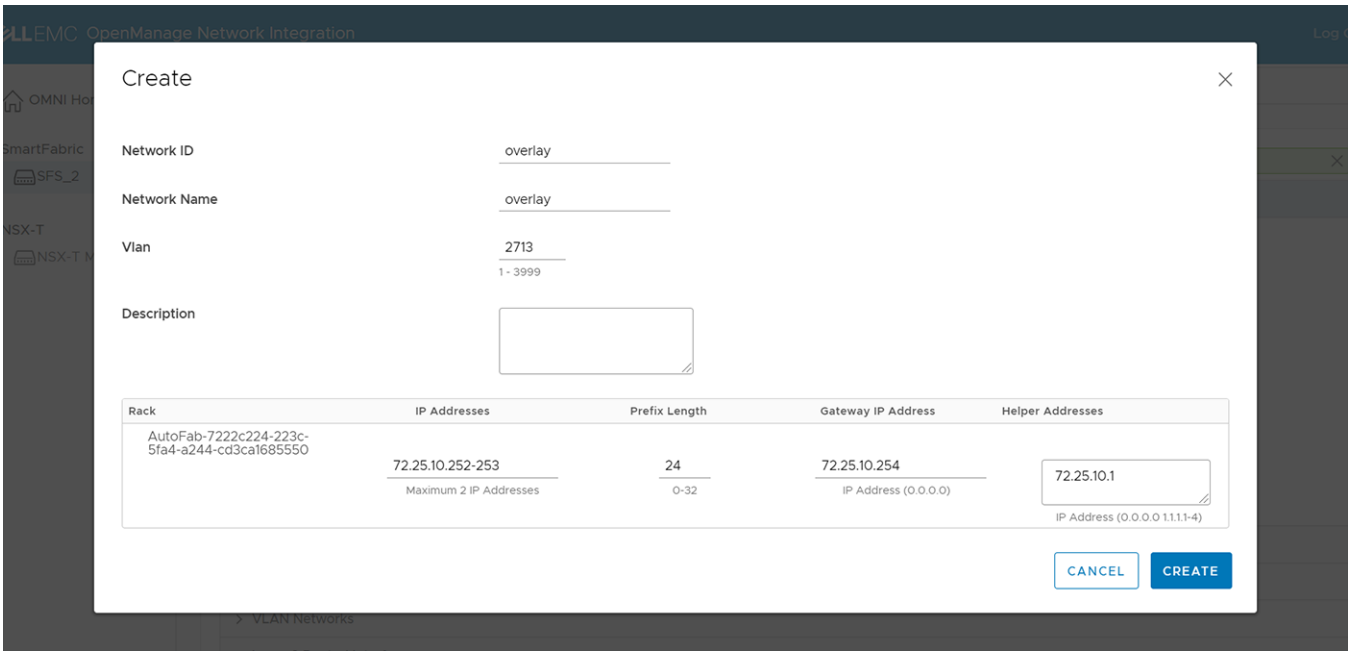

3. The system displays VLAN network creation success message.

Edit multi rack L3 VLAN configuration

1. Select a network ID from the list, and click **Edit**.

2. Modify the details, edit the configuration as necessary, and click **Edit**.

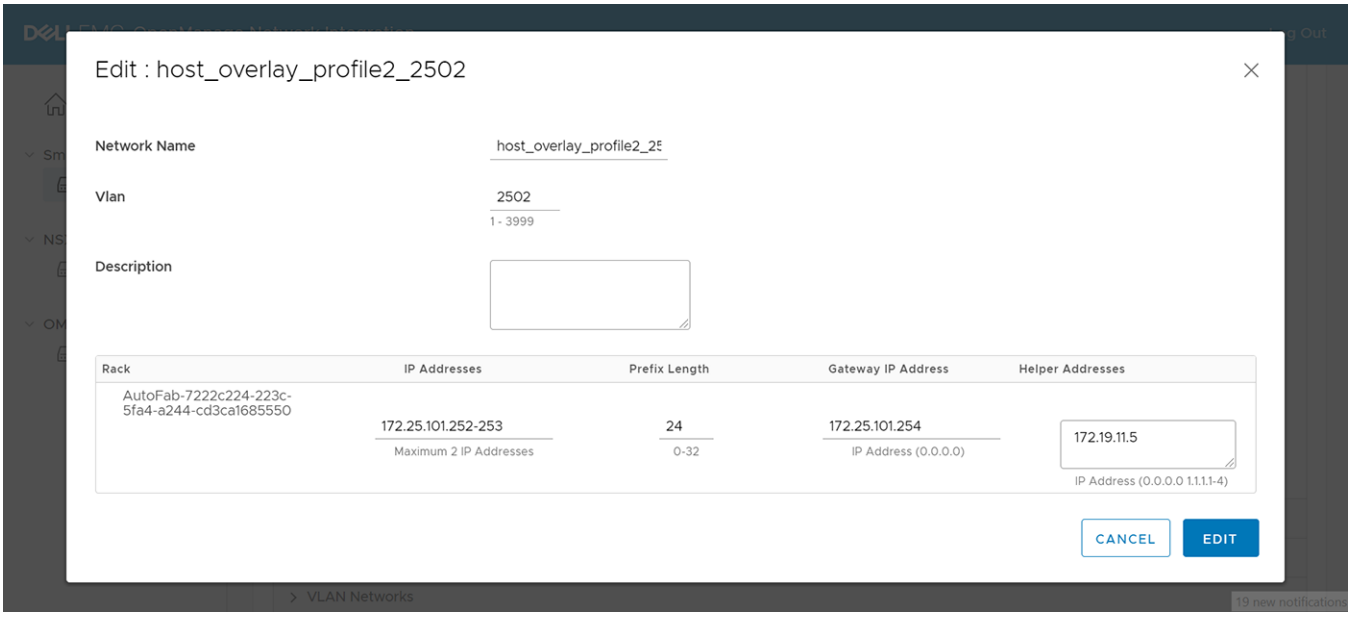

3. The system displays edit network success message.

Delete multi rack L3 VLAN configuration

1. Select the VLAN network to remove, and click **Delete**.

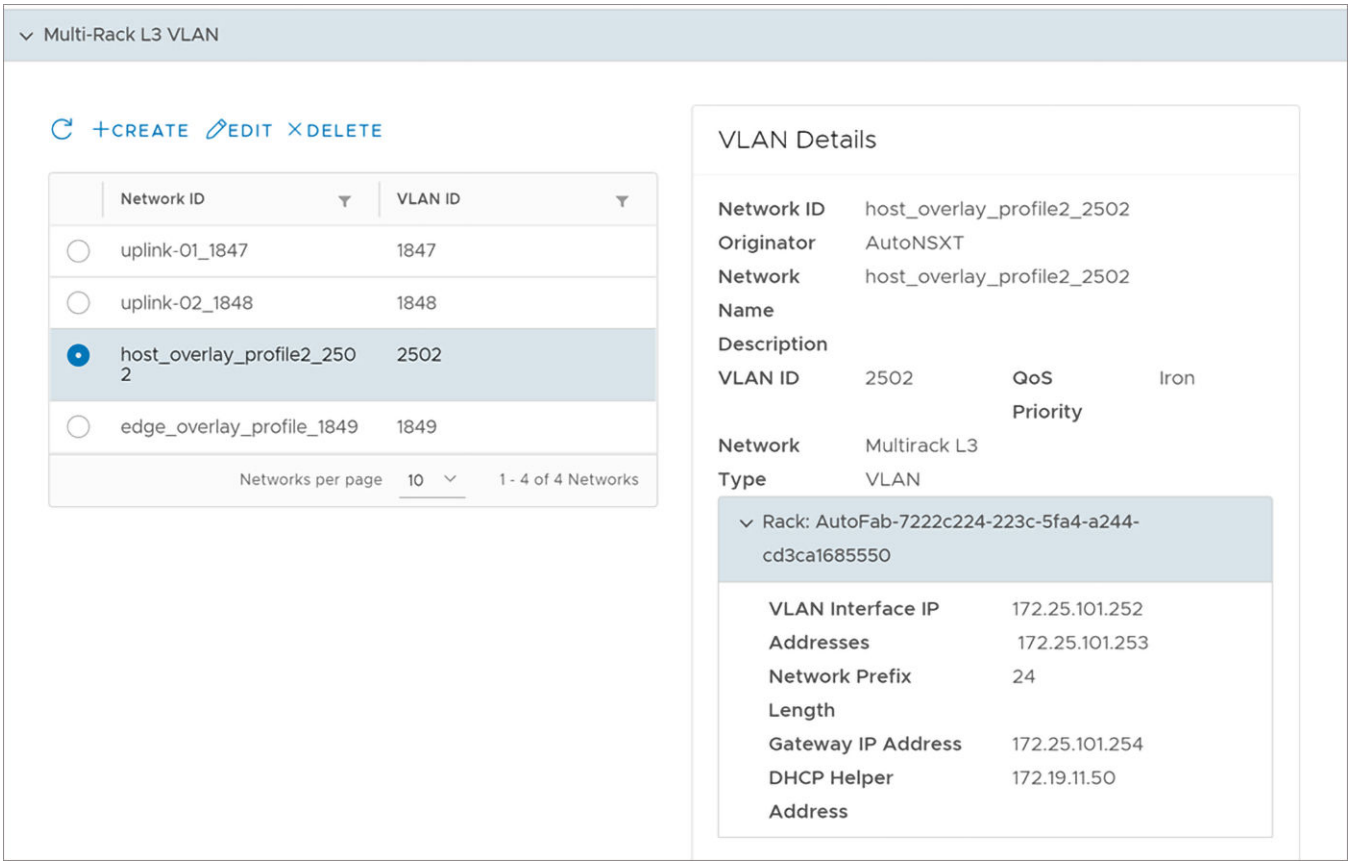

2. Click **Delete** to confirm.

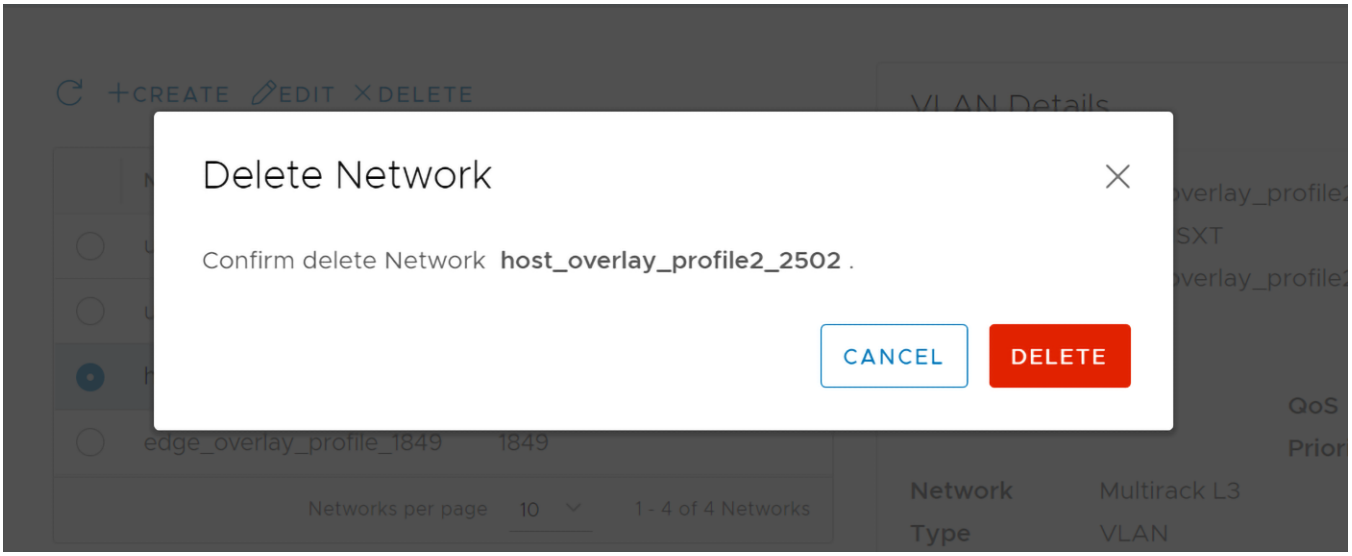

3. The system displays network deletion success message.

Configure general purpose networks

When you create a general purpose network, OMNI 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, OMNI 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.

NOTE: OMNI UI does not display the virtual networks that are created automatically during general purpose network creation, as OMNI UI is designed to filter these virtual networks when displayed in the UI. However, SFS UI displays the virtual networks that are created automatically during the general purpose network creation.

Create general purpose network

To create a general purpose network:

1. Click the SFS instance for which you want to create a network.
2. Click **Networks** > **General Purpose Networks**. The page displays the list of the general purpose networks that are already configured in the SmartFabric.
3. Click **Create** to create a Layer 2 general purpose network.
4. Enter the following details:
   - Network ID.
   - Network name. For example, network-201.
   - VLAN ID. A number that ranges from 1 to 3999 (except 3939). For example, 201.
   - Description.

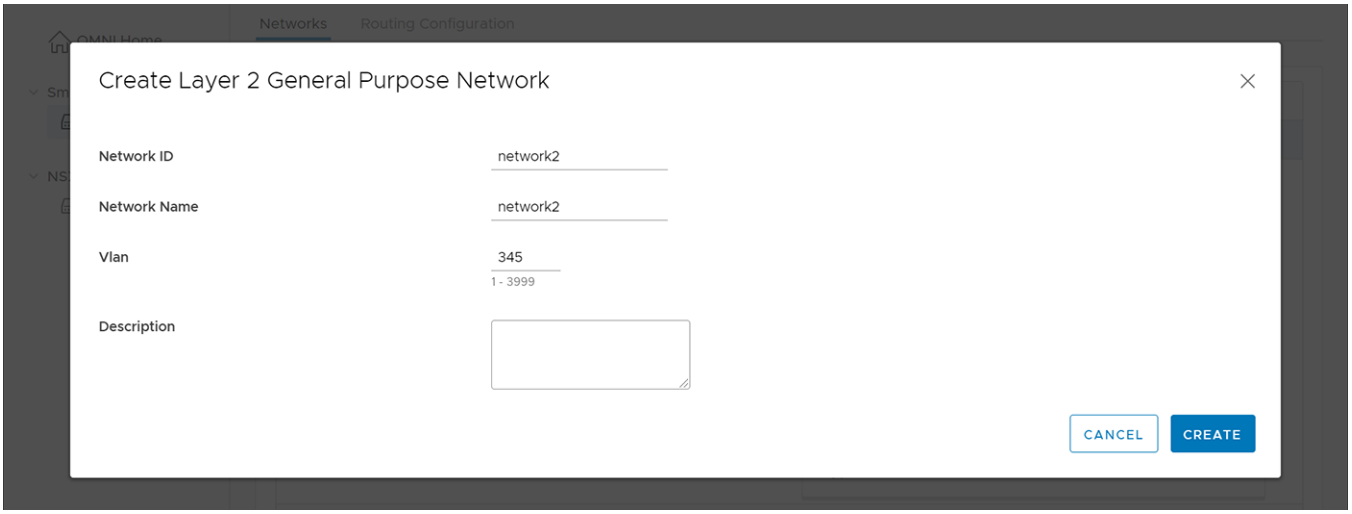

5. Click **Create**. The system displays virtual network creation successful message.

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**.

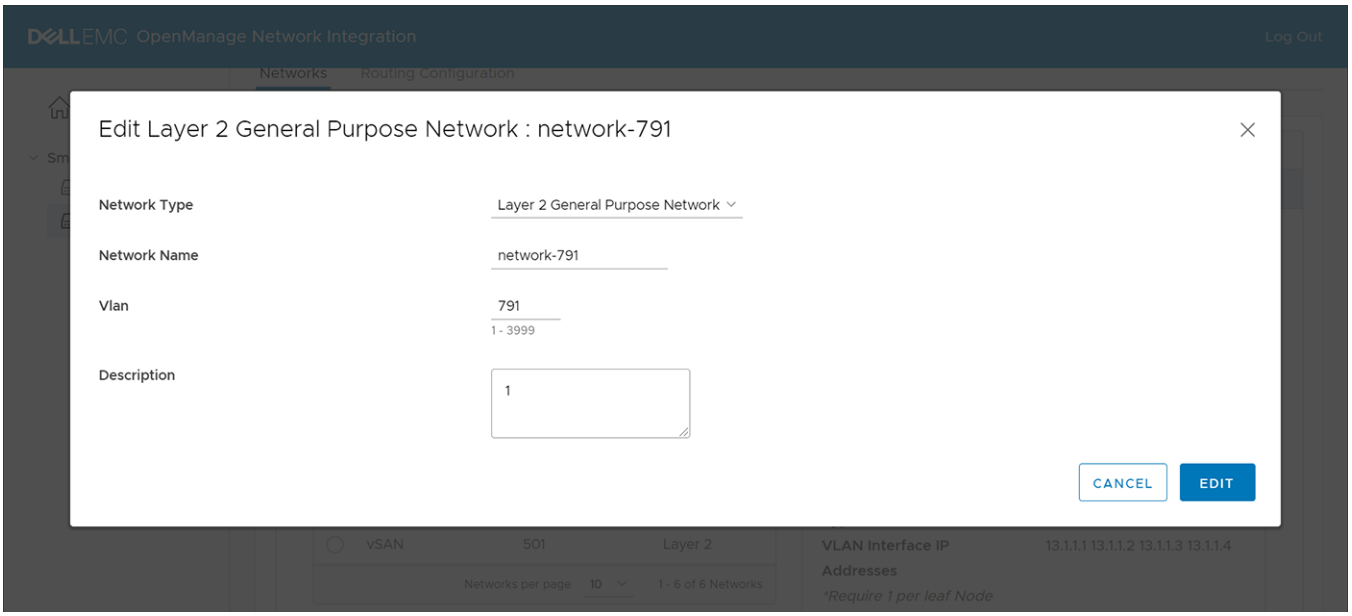

2. Select the network type to Layer 3 general purpose network.
3. Enter the following details:
   - Network name
   - VLAN
   - Description
   - IP address
   - Prefix length
   - Gateway IP address
   - Helper address

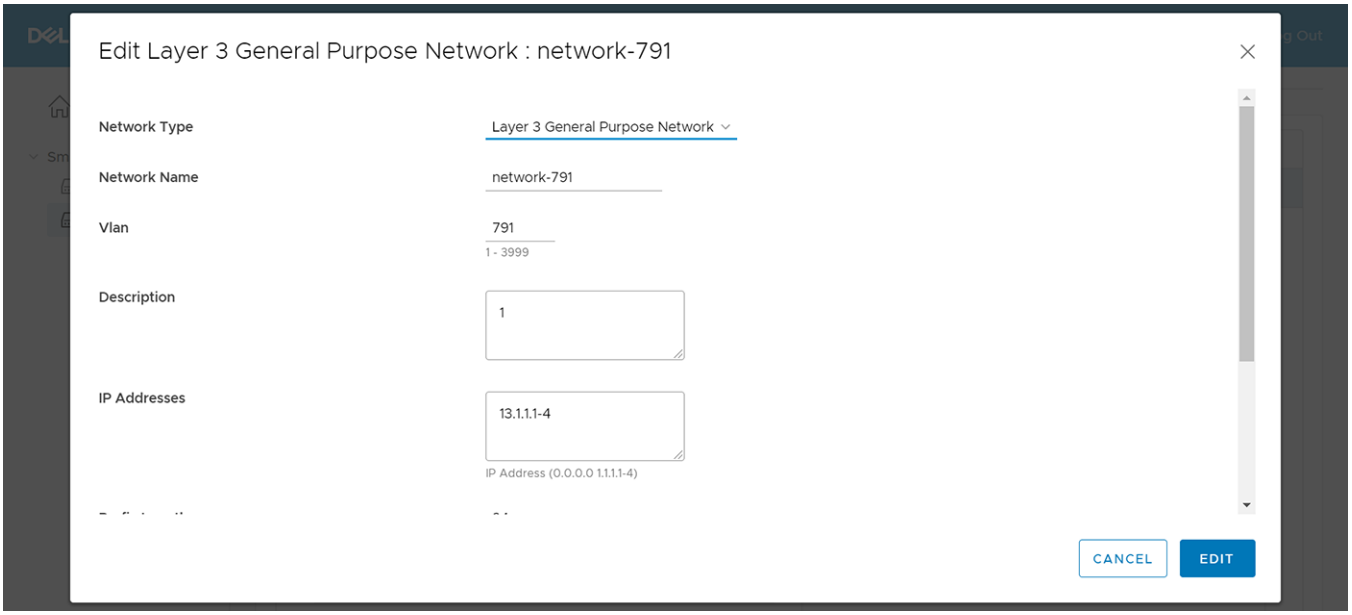

4. Click **Edit**. The system displays virtual network edits success message.

View general purpose network

To view the details of the general purpose networks, select a network from the list. The VLAN details of the specific network including network ID, originator, network name, VLAN ID, QoS priority, network type, VLAN interface IP address details, prefix length, gateway IP address, and DHCL helper address. `Portgroups` that are created on the vCenter are displayed under **General Purpose Networks**.

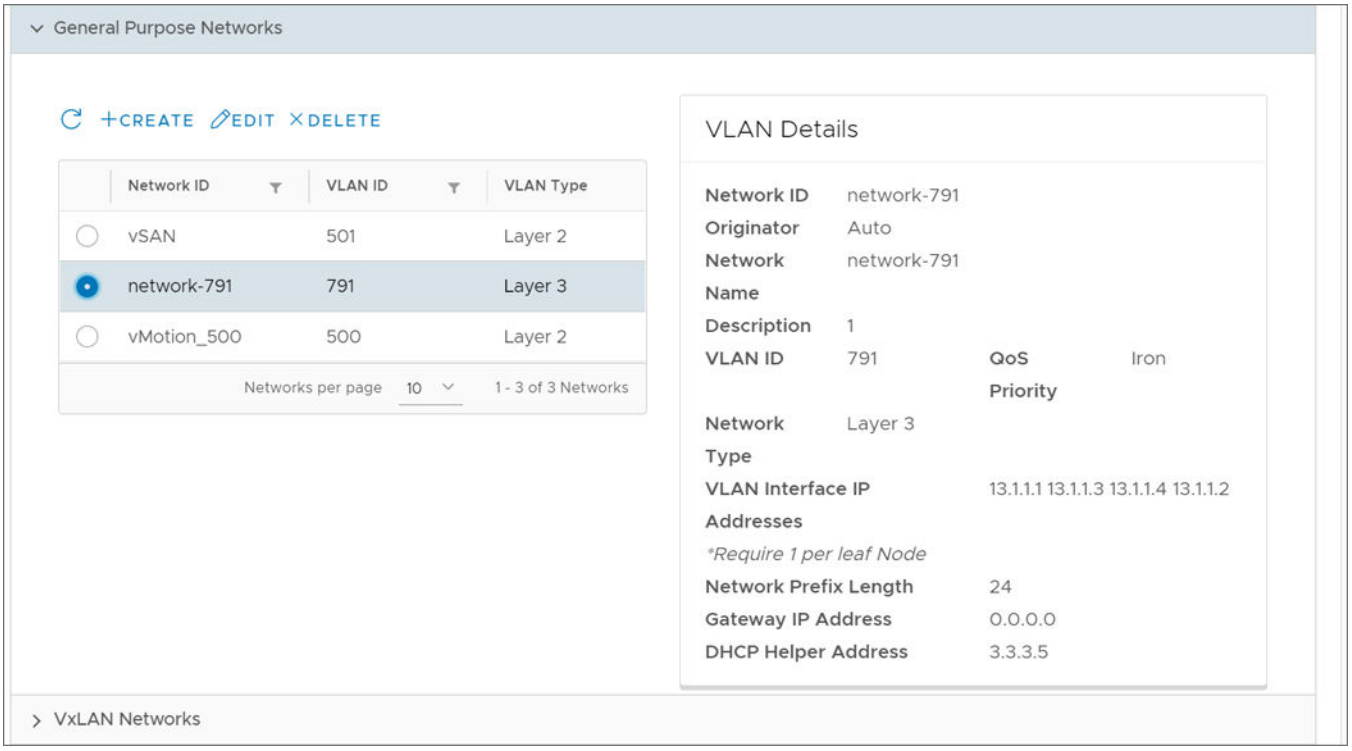

Delete general purpose network

When you delete a general purpose network, both the VLAN and the VXLAN networks are deleted from OMNI. 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 VxLAN network

Create, edit, and delete L2 and L3 profile VXLAN network configurations through OMNI. 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. From SmartFabric instance, select **Networks** > **VxLAN Networks**. The page displays the list of the VXLAN networks that are configured in the SmartFabric instance.

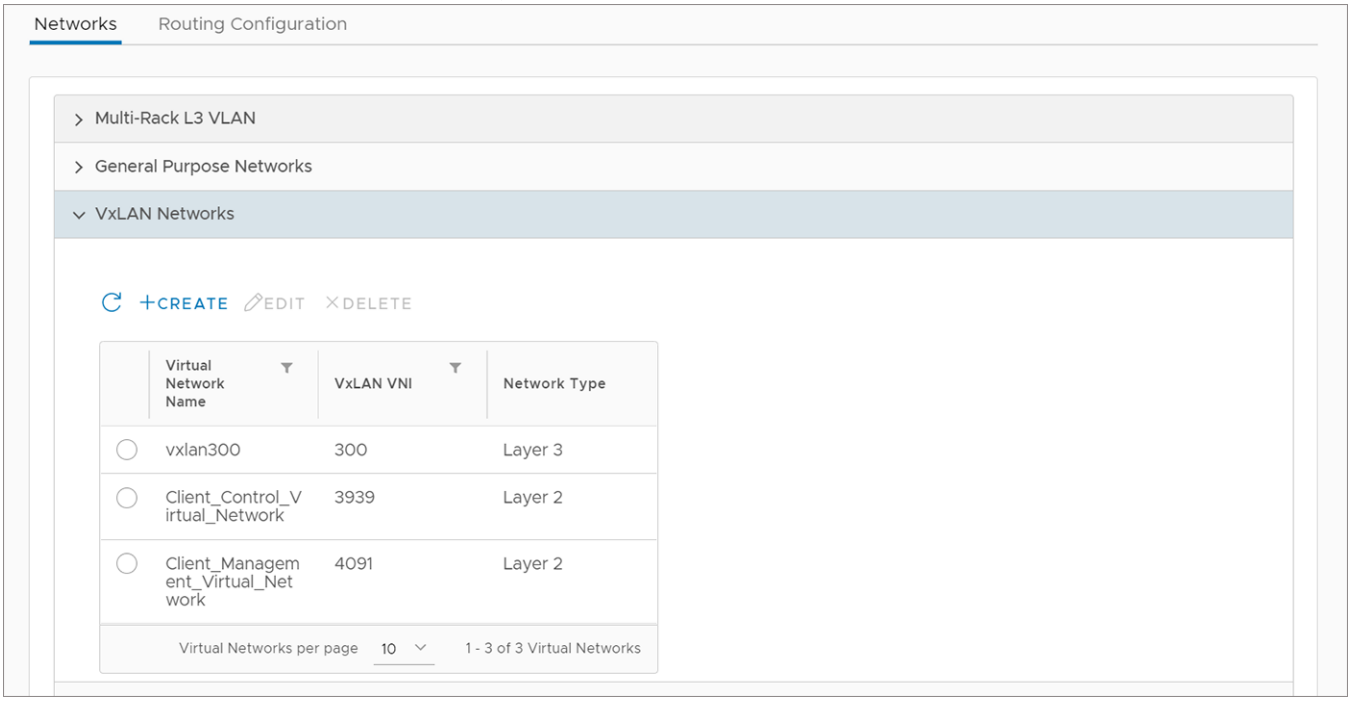

2. Click **Create**.
3. Verify **Layer 2** is selected as the **Virtual Network Type**.
4. Enter the text for **Virtual Network Name**, a value for the VxLAN VNI, and the VLT VLAN ID.

5. (Optional) Enter a description, and click **Create**.

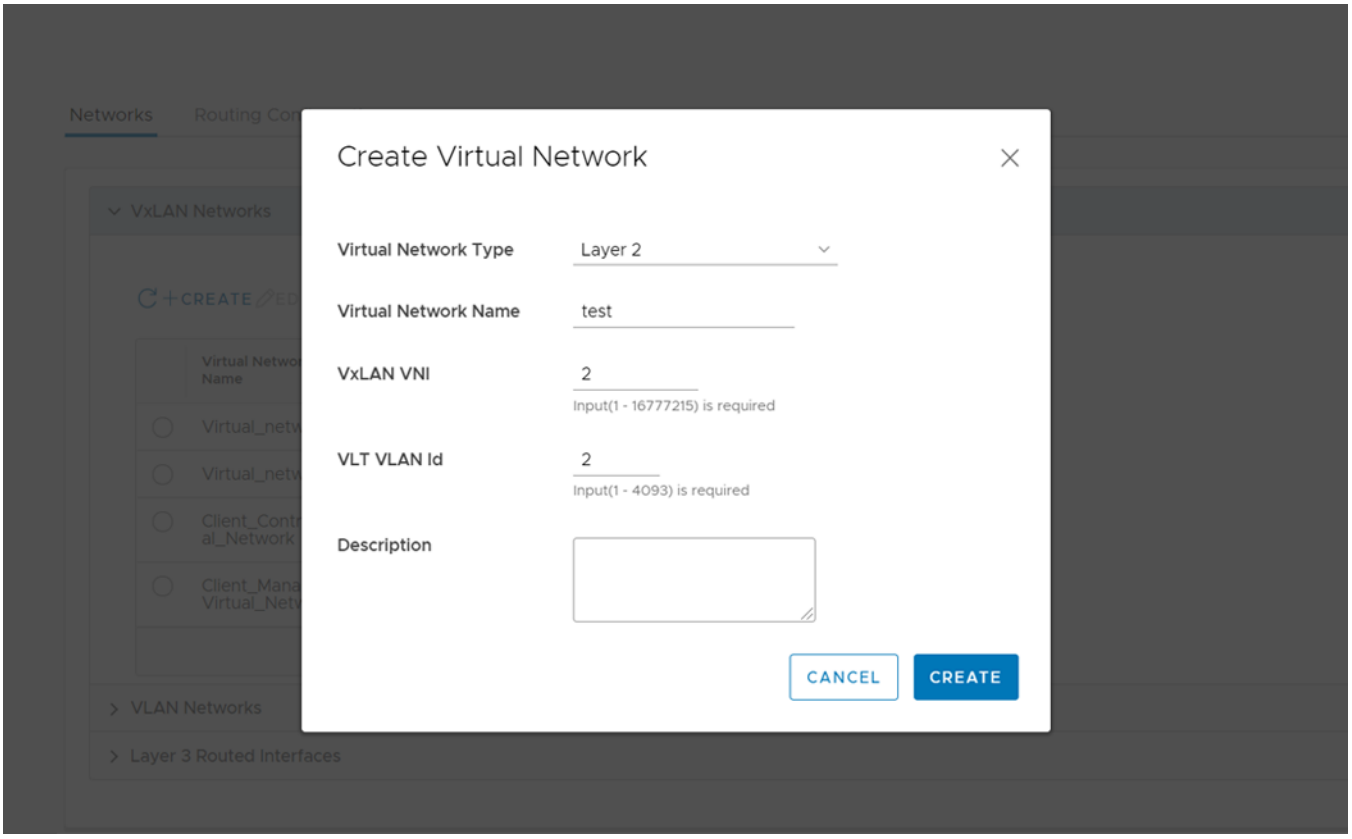

6. The system displays virtual network creation successful message.

Virtual network for L3 profile:

1. Select Network from the Network tab, then click **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 text for **Virtual Network Name**, a value for the VxLAN VNI, the VLT VLAN ID, prefix, gateway IP address, and helper IP address. Click **Create**.

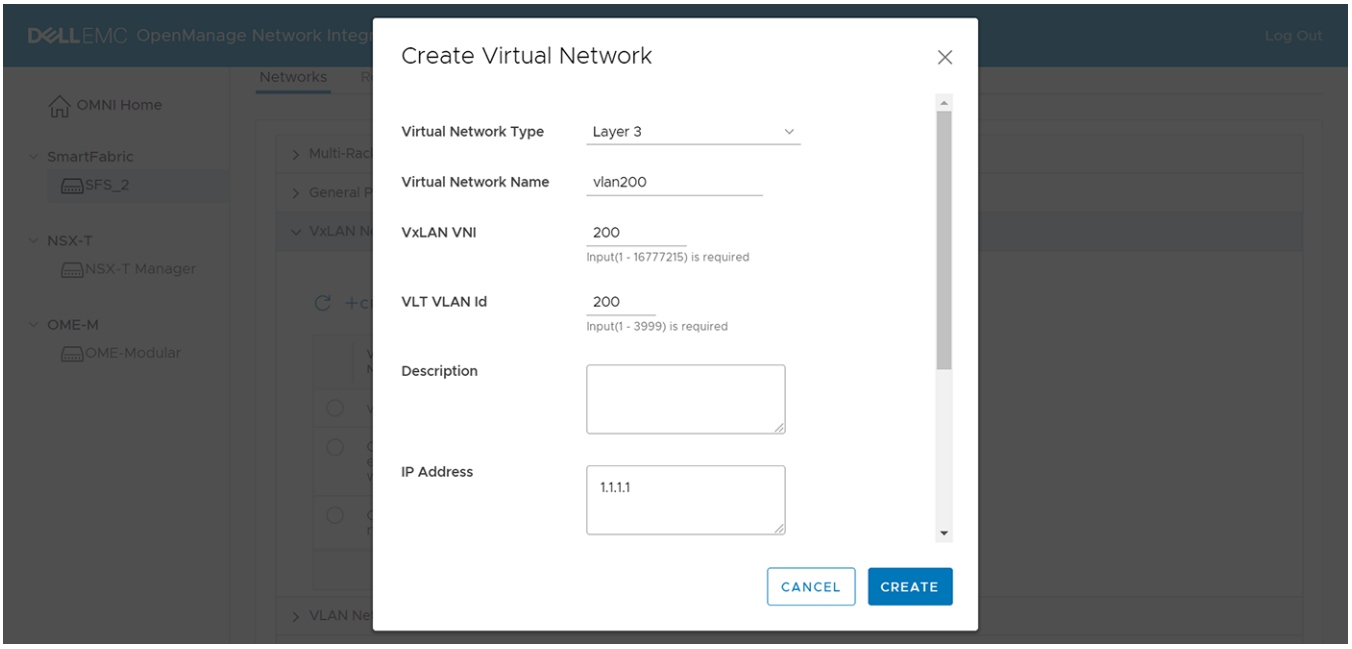

5. The system displays virtual network creation successful message.

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.

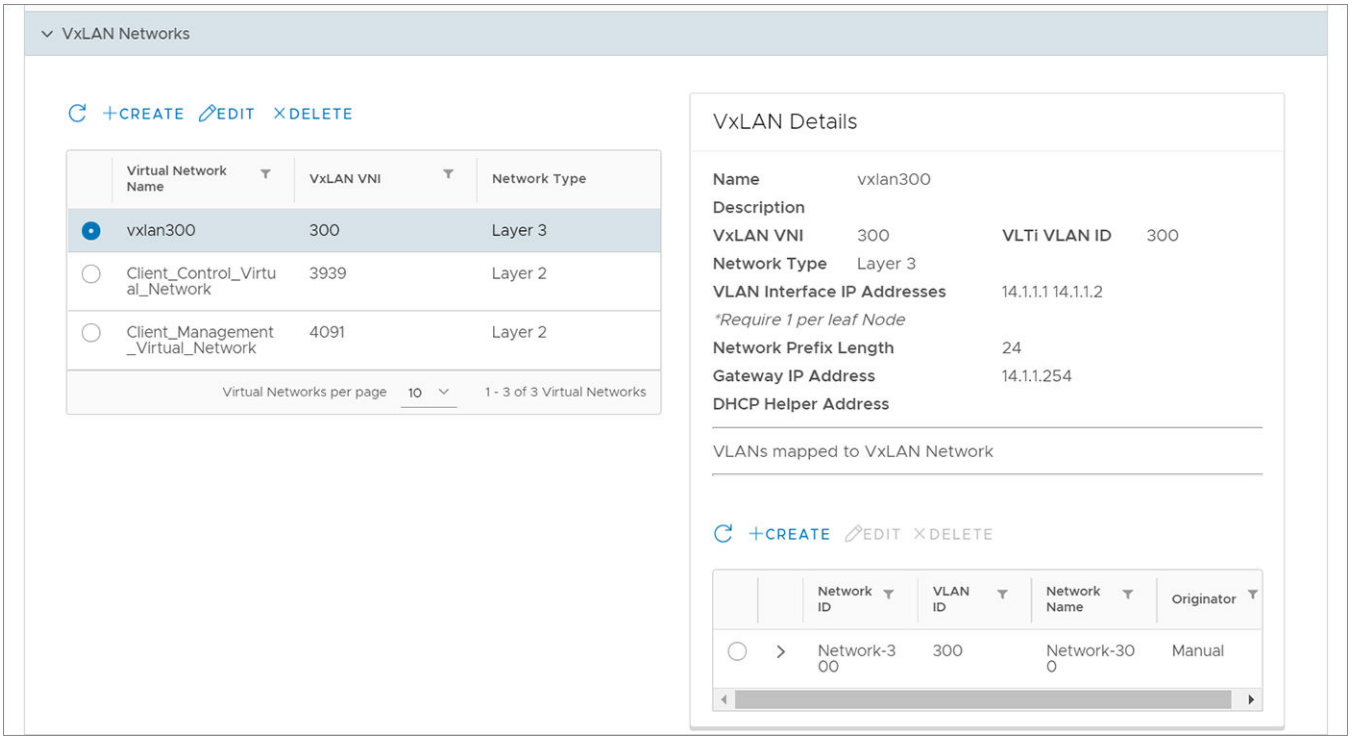

Edit VxLAN network

You can edit the configuration of VXLAN network:

1. Select a virtual network from the list, then click **Edit**.

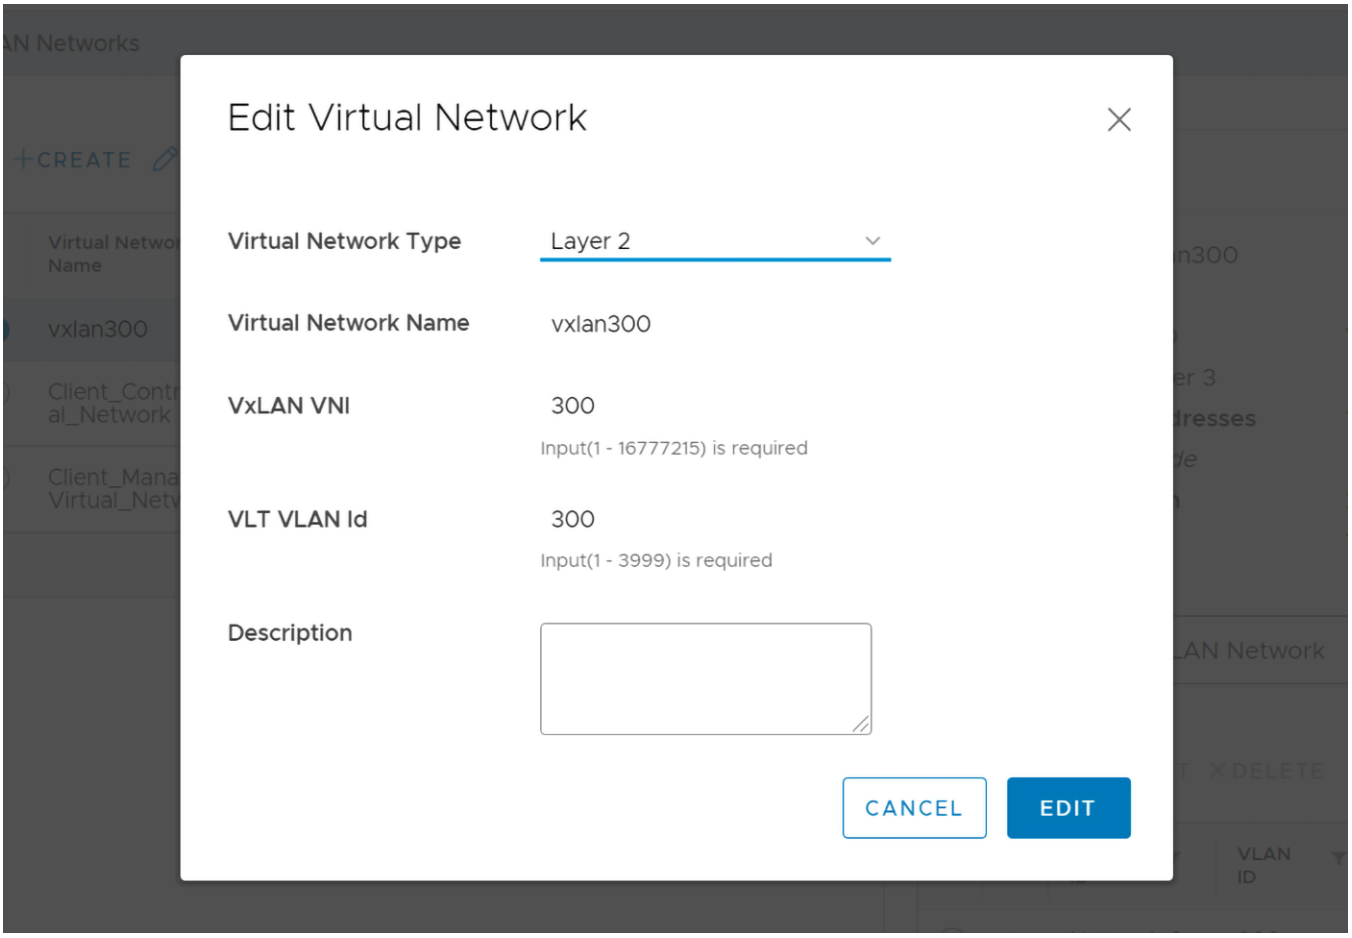

2. Modify the Virtual Network Type.

3. Enter the Prefix, Gateway IP Address, IP address, then click **Edit**.

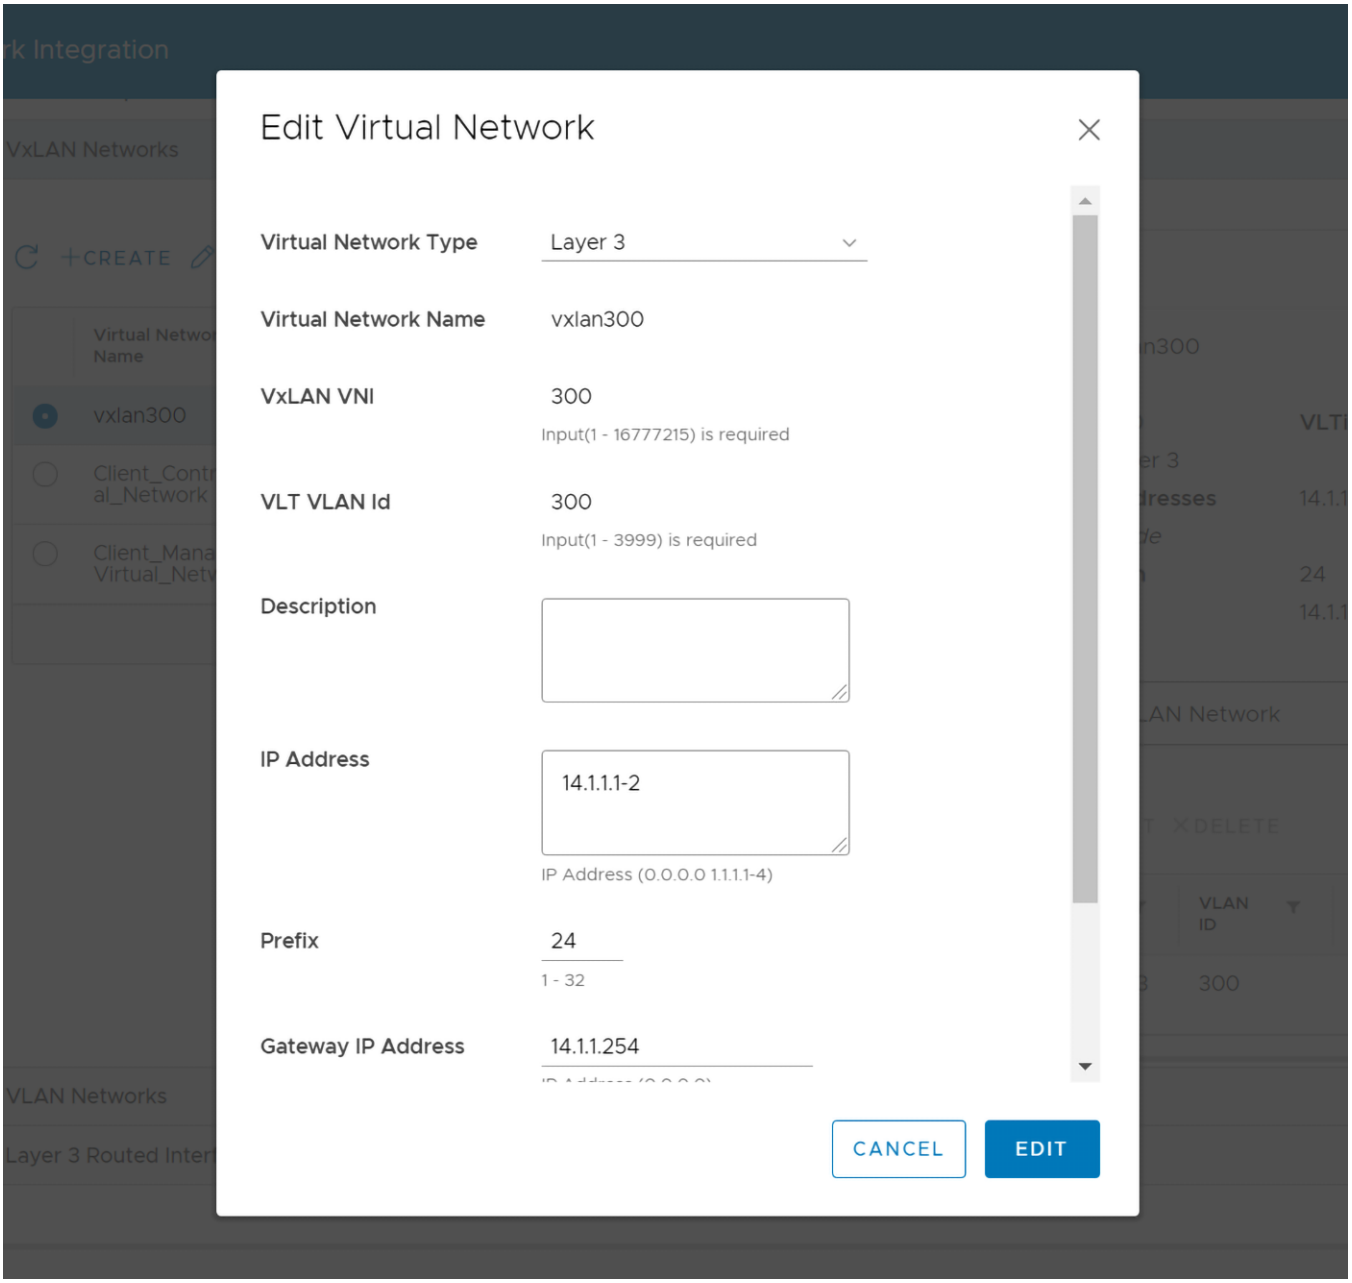

4. The system displays virtual network edits success message.

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**.

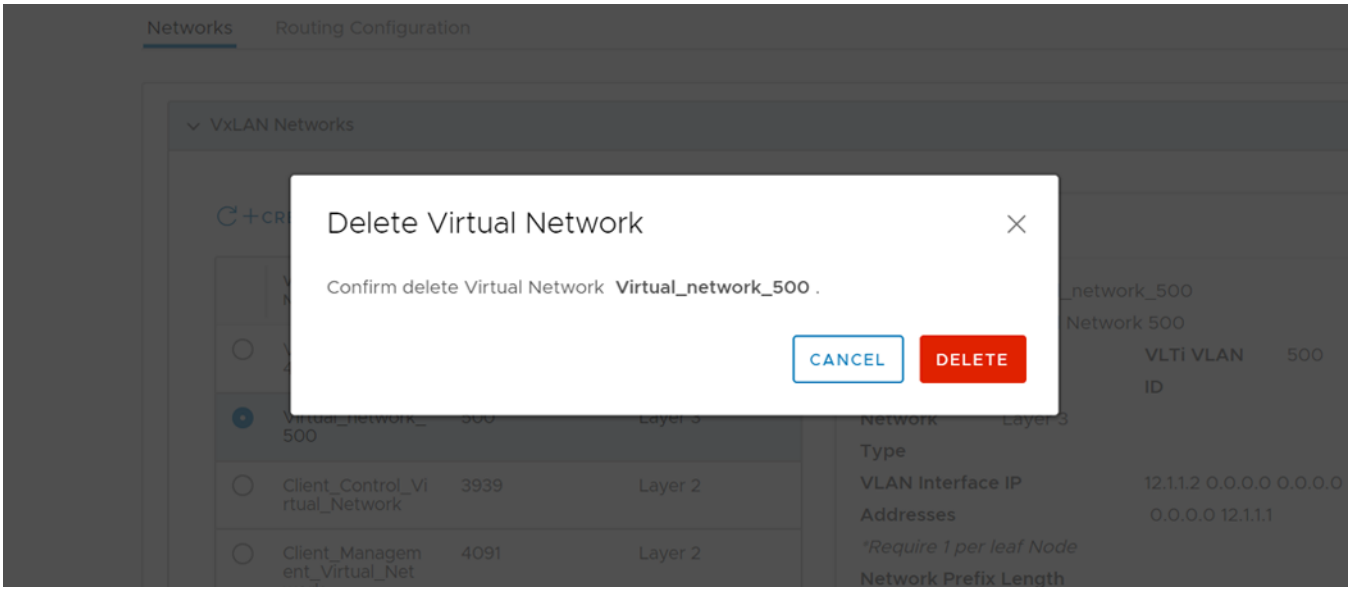

2. Click **Delete** to confirm.
3. The system displays network deletion success message.

Configure VLAN networks

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

Create L2 VLAN or L3 VLAN network

VLAN networks for L2 profile:

1. Select **Networks** > **VLAN Networks**, and click **Create**.

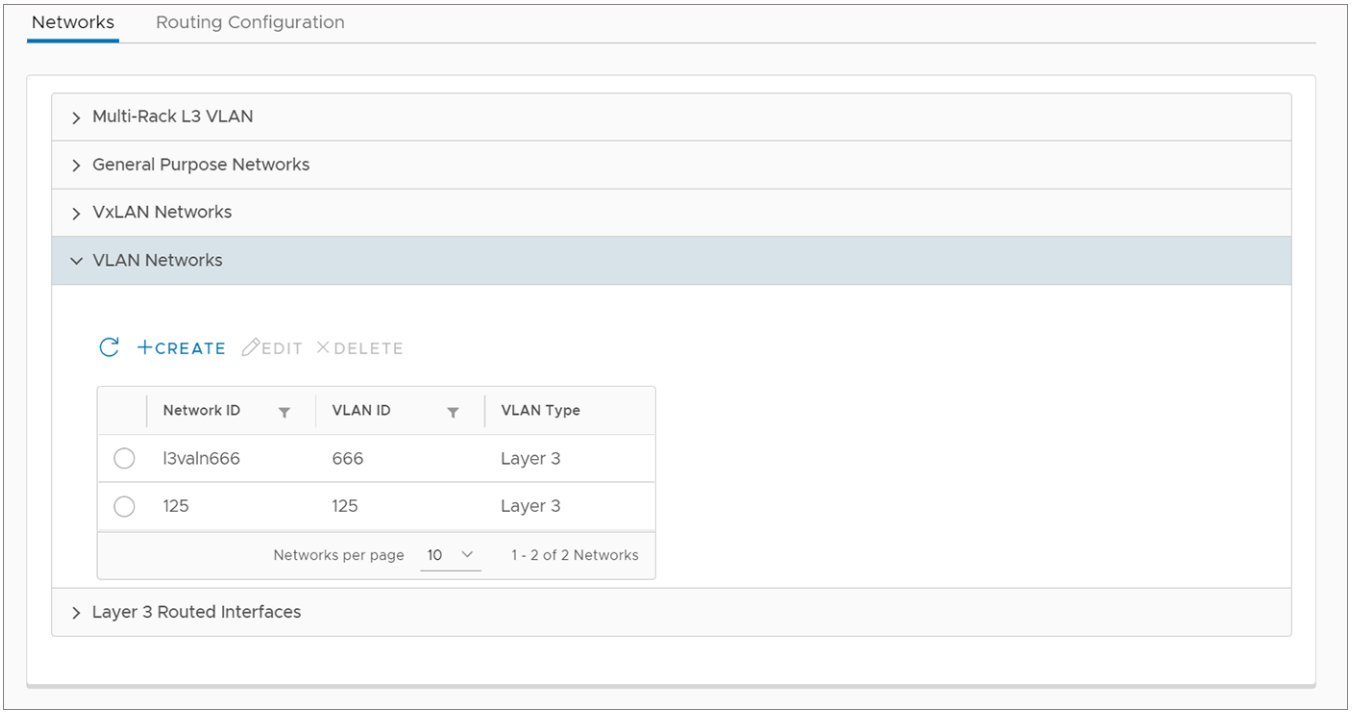

2. Select the Network Type as **Layer 2 VLAN Network** is selected as the Network Type, enter the **Network ID**, **Network Name**, enter a number for the VLAN, enter an optional description, then click **Create**.

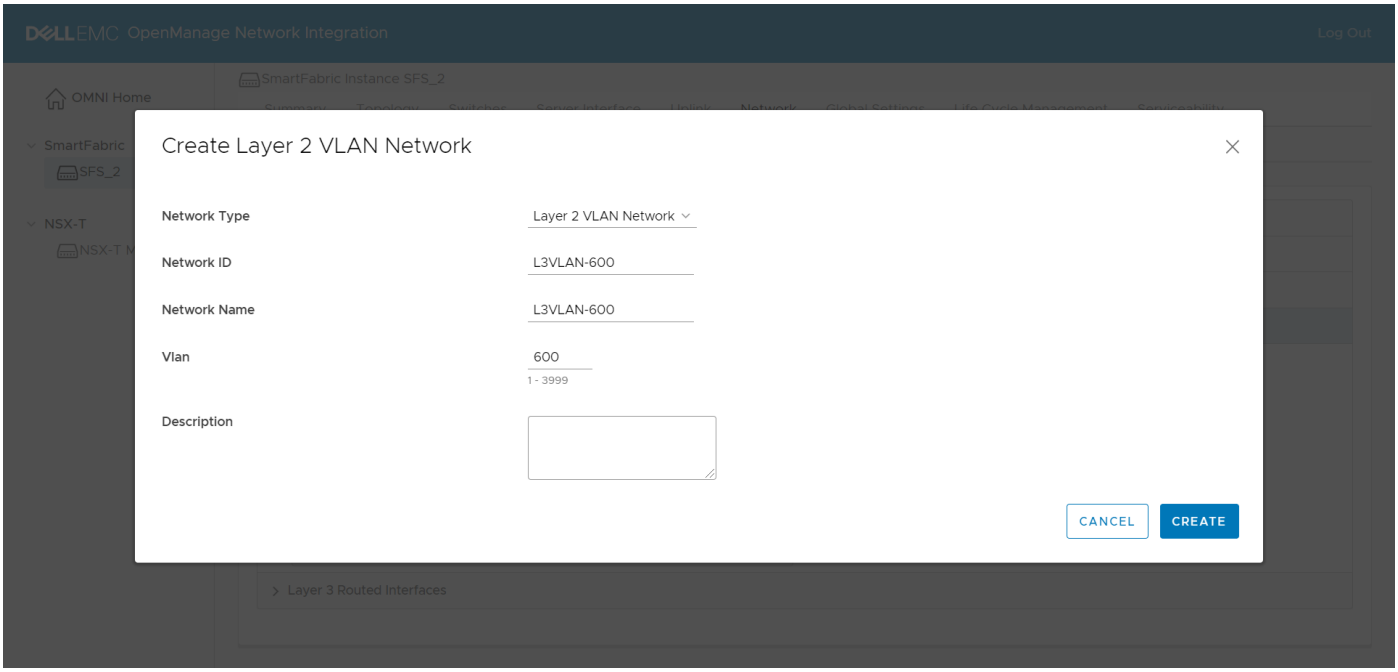

3. The system displays VLAN network creation success message.

VLAN networks for L3 profile:

1. Select **Networks** > **VLAN Networks**, and click **Create**.
2. Select the Network Type as **Layer 3 VLAN Network** is selected as the Network Type, enter the **Network ID**, **Network Name**, enter a number for the VLAN, enter an optional description, then click **Create**.

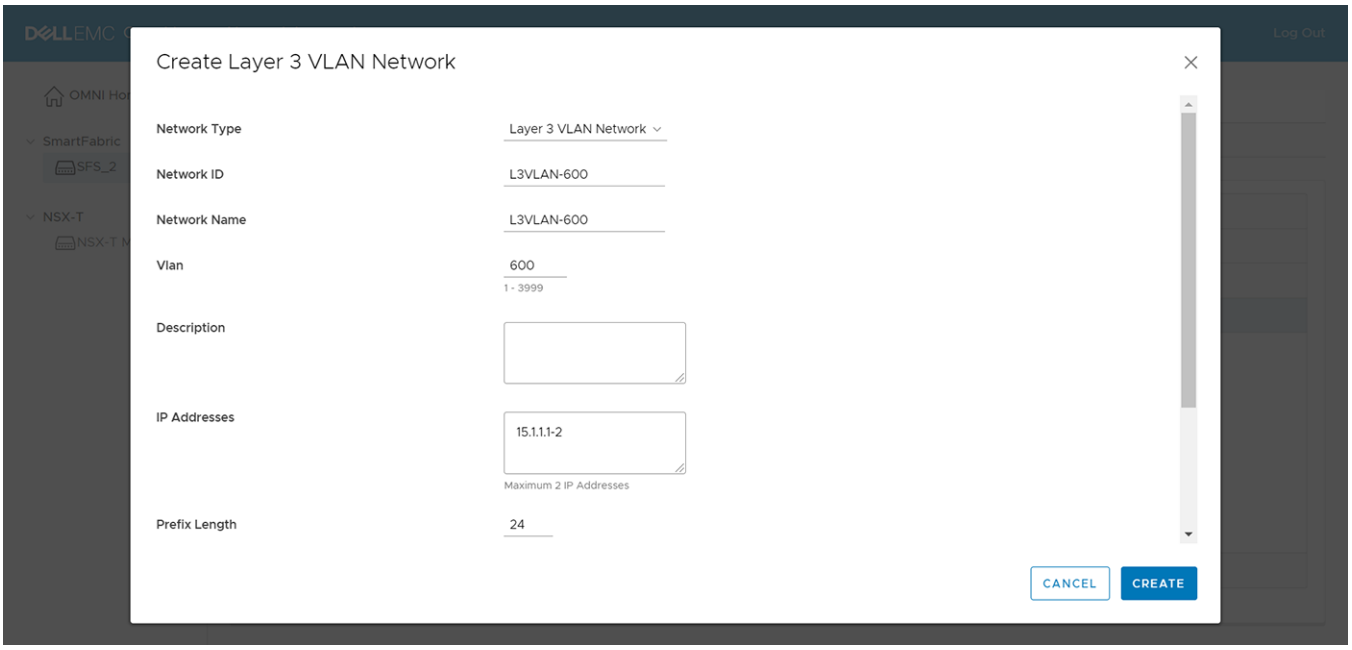

3. The system displays VLAN network creation success message.

Edit network

1. Select a network ID from the list, and click **Edit**.

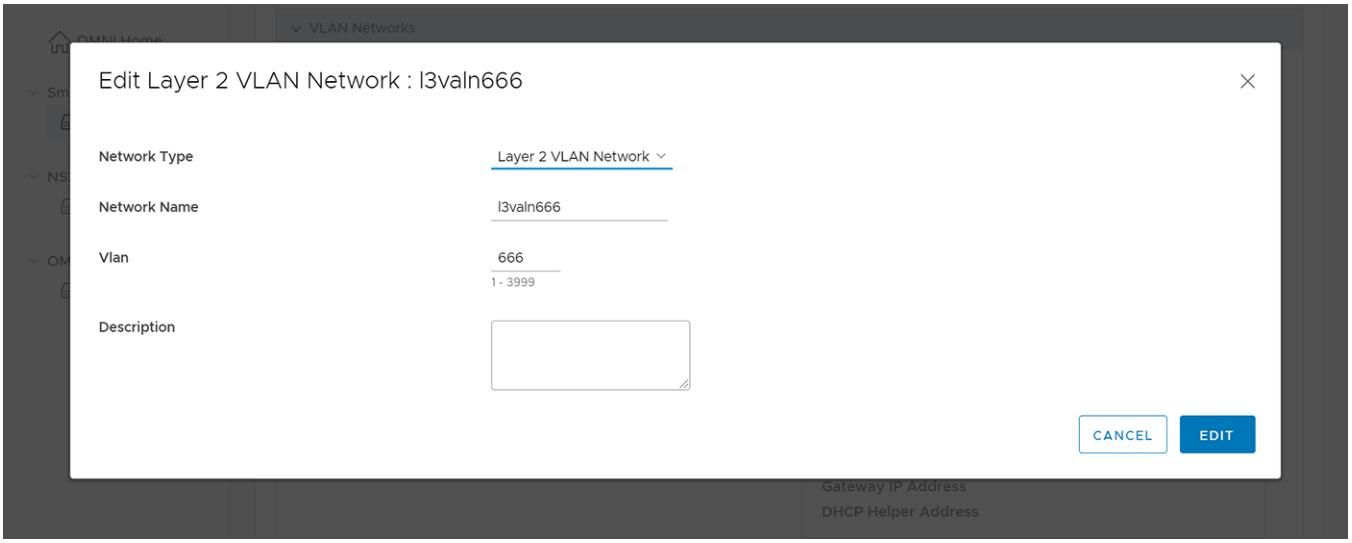

2. Modify the details, edit the configuration as necessary, and click **Edit**.

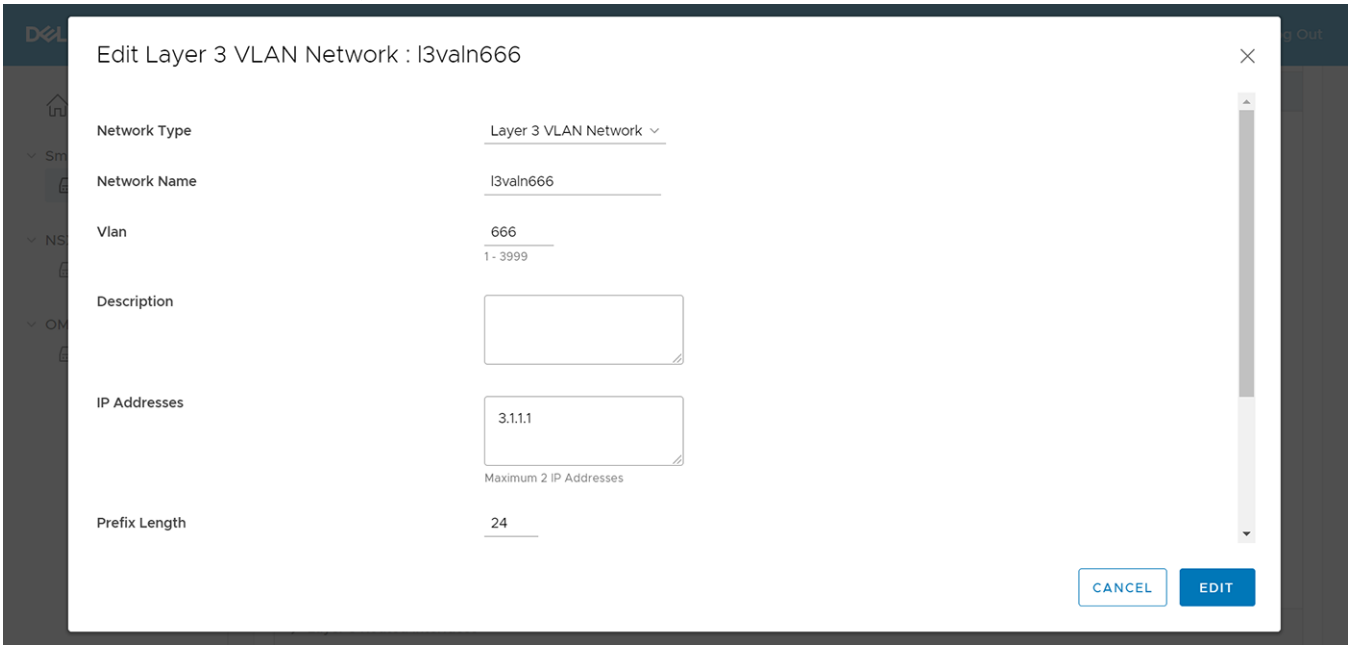

3. The system displays edit network success message.

Delete network

1. Select the VLAN network to remove, and click **Delete**.

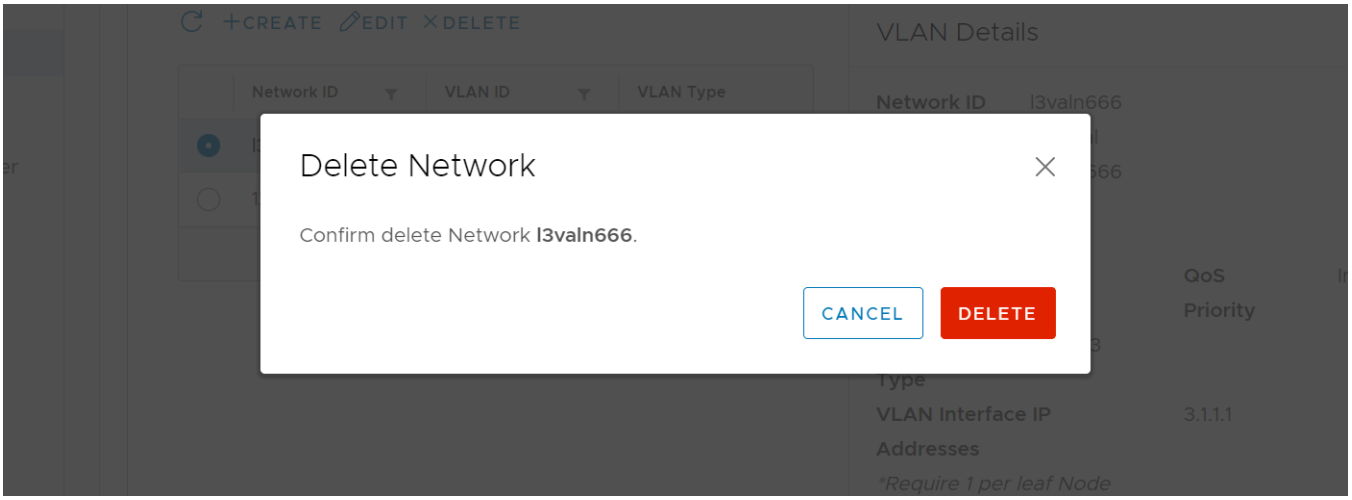

2. Click **Delete** to confirm.
3. The system displays network deletion success message.

Configure L3 routed interfaces

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

Create L3 routed interface

To create an L3 routed interface:

1. Select **Networks** > **Layer 3 Routed Interfaces**, and click **Create**.

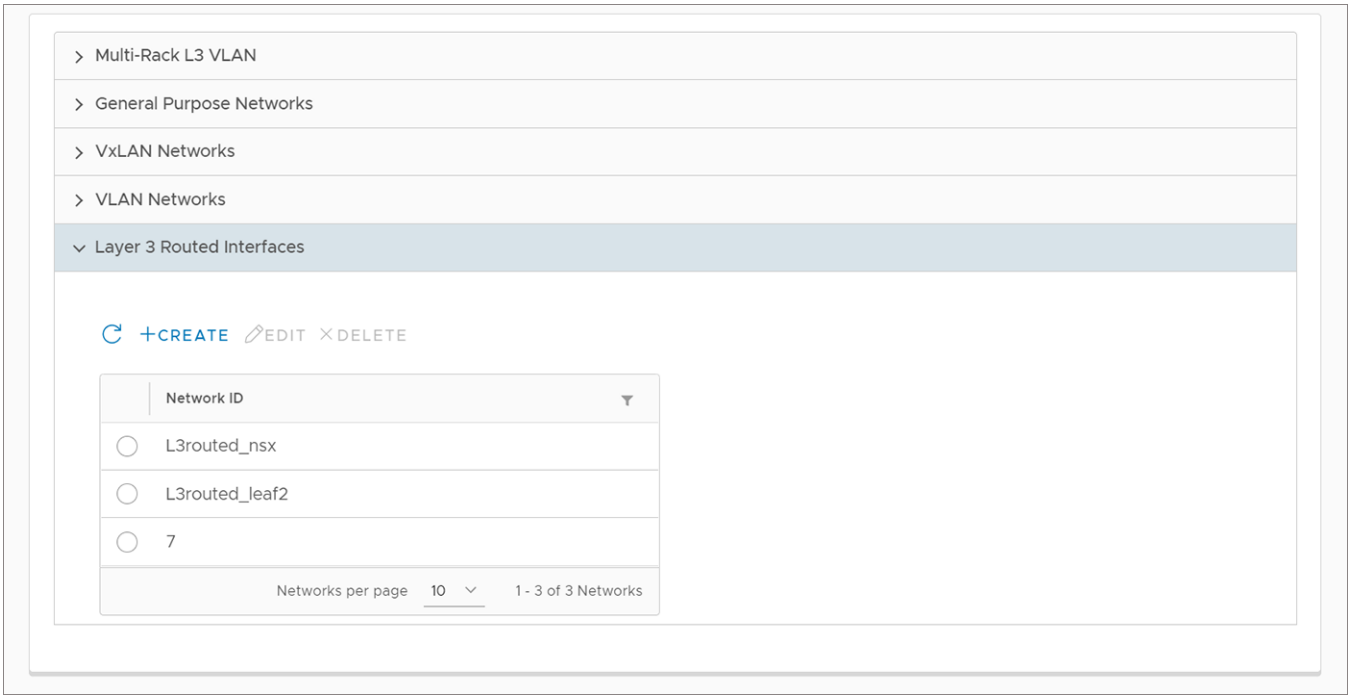

2. Enter the **Network ID**, **Network Name**, select the **Prefix Length**, select the **IP Address**, enter an optional description, then click **Create**.

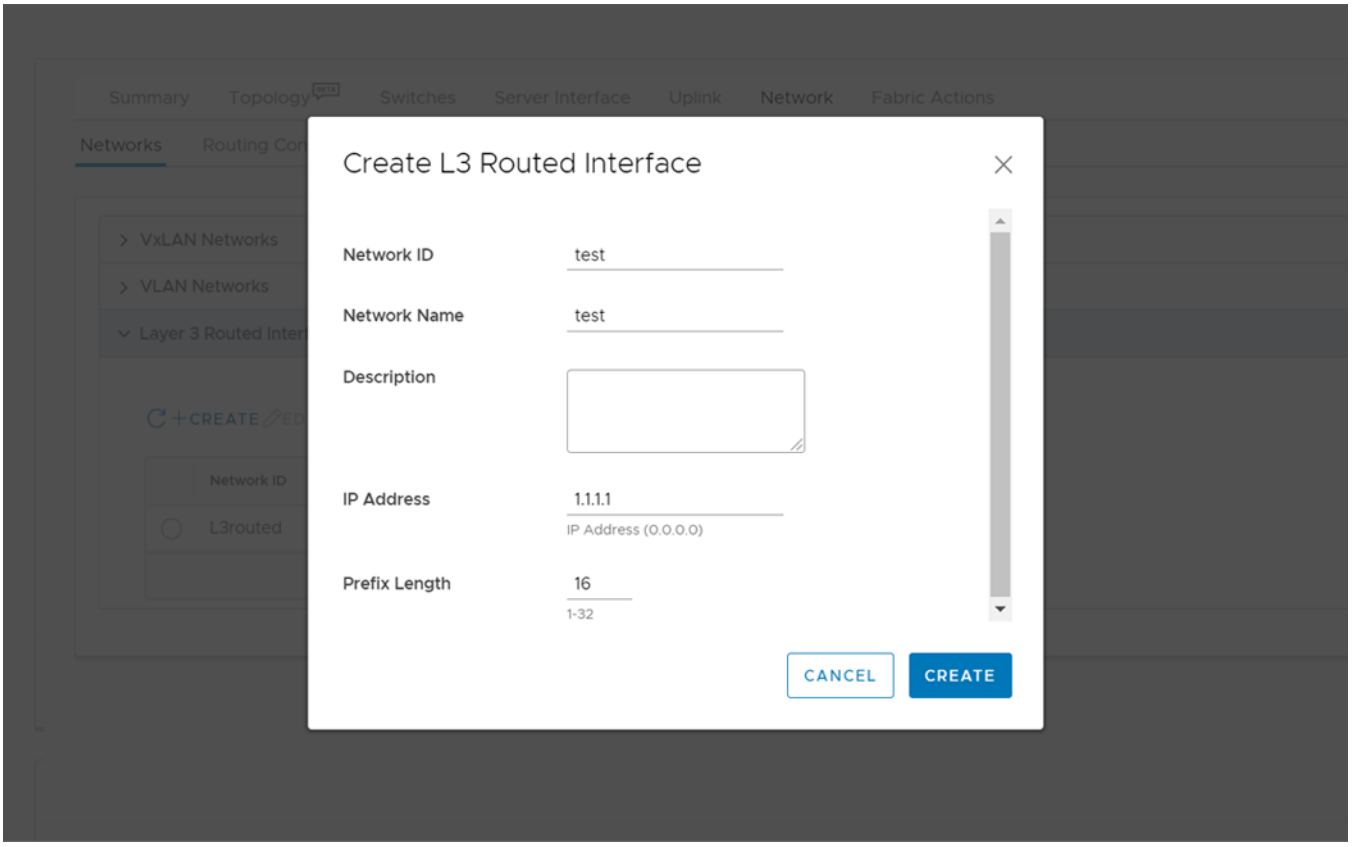

3. The system displays network creation success message.

Edit network

1. Select the **Network ID** from the list, and click **Edit**.

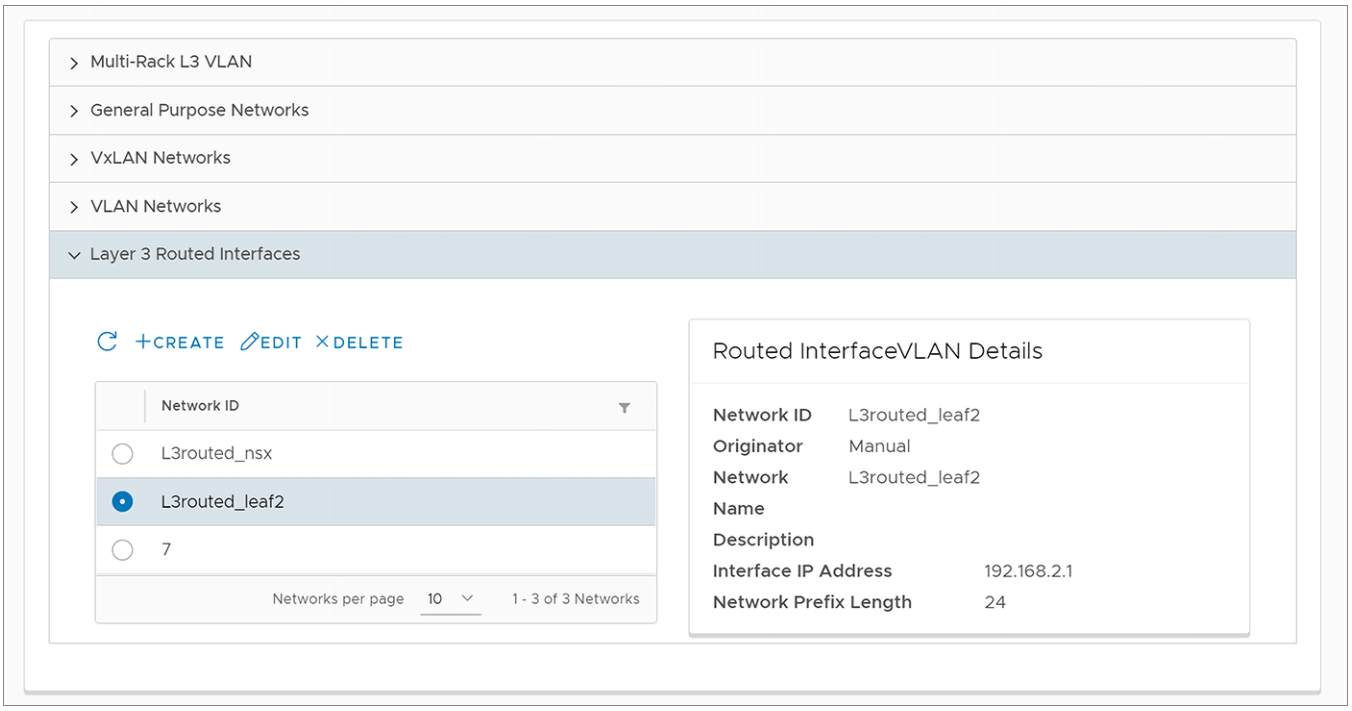

2. Edit the configuration, and click **Edit**.

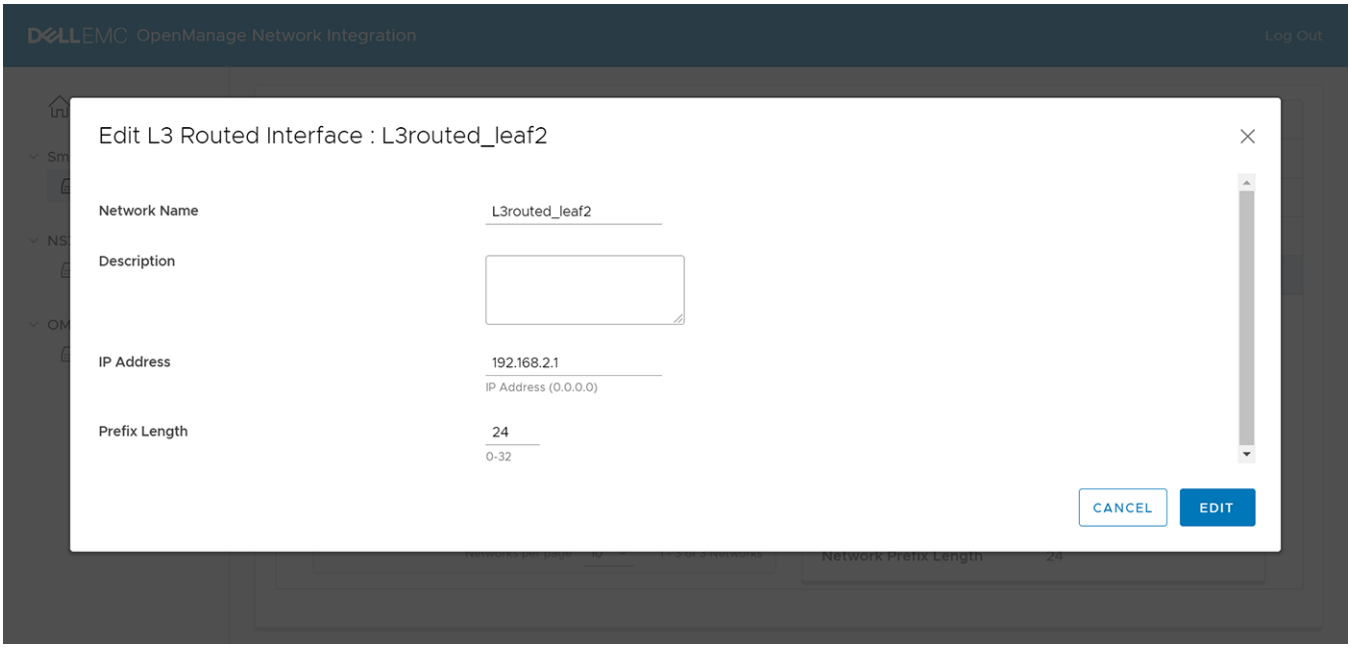

3. The system displays edit network success message.

Delete network

1. Select the network ID to remove, and click **Delete**.

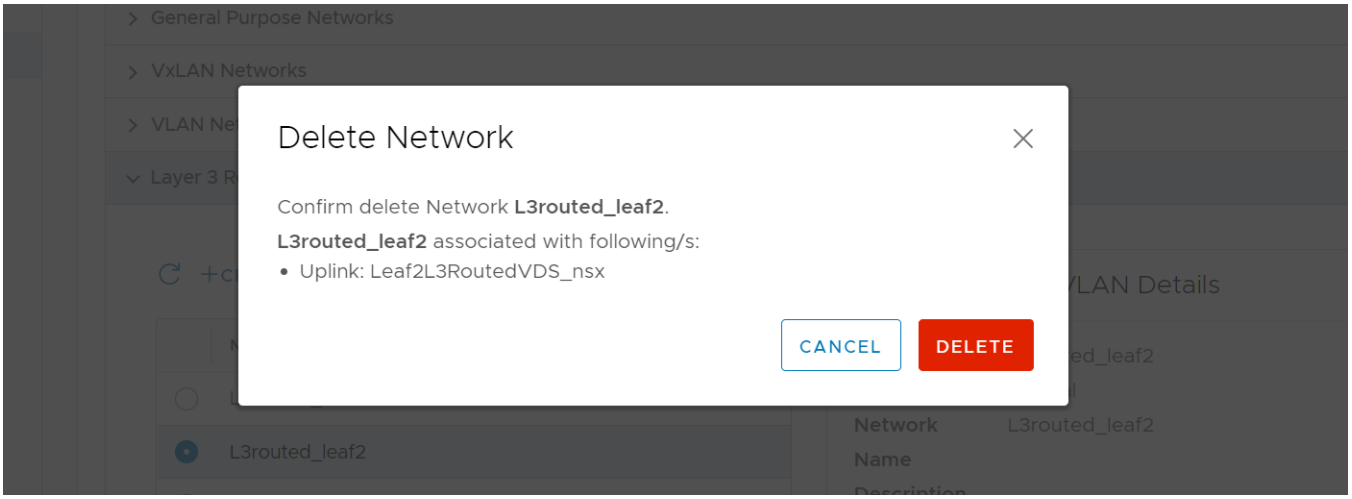

2. The system displays network deletion success message.

Configure Routes

You can configure static routes and eBGP peer routes for a network.

Configure static routes

Create static route

1. Select **Network** > **Routing Configuration**.

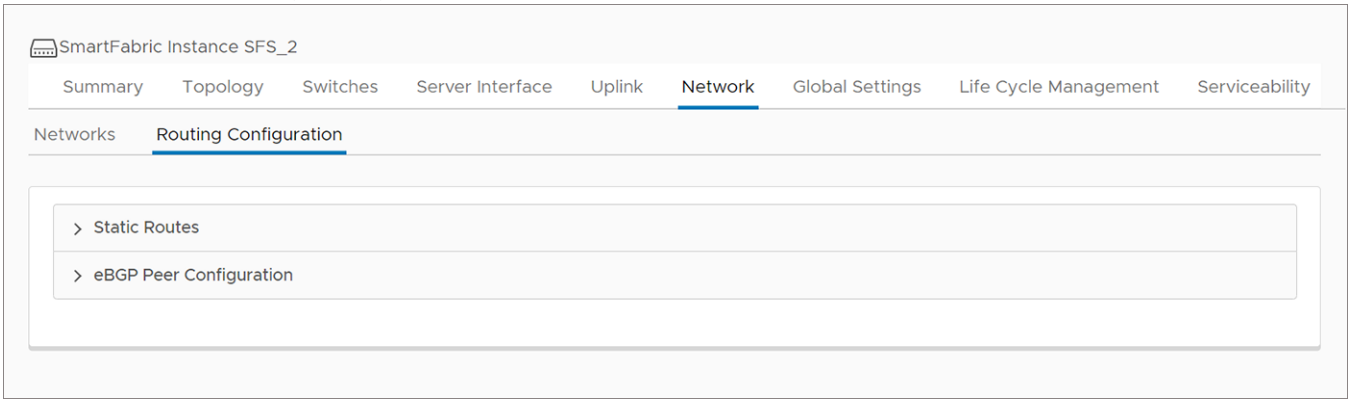

2. Select **Static Routes**, and click **Create** to add a new static route.
3. Enter the relevant details and click **Create**.

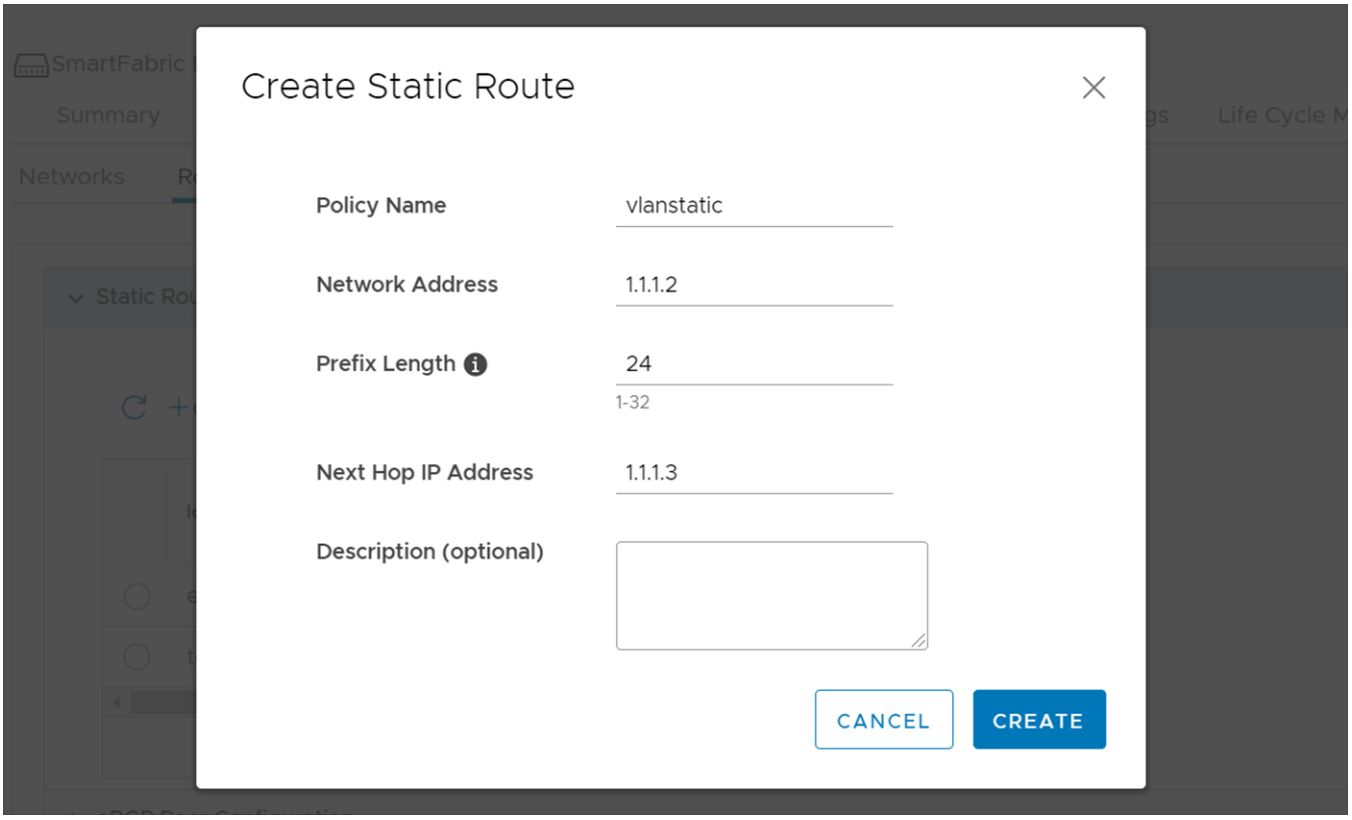

4. The system displays static route creation is successful.

Add route to switch

1. Select **Routing Configuration** > **Static Routes**.

2. Select a static route, and click **Add Route to Switch**.

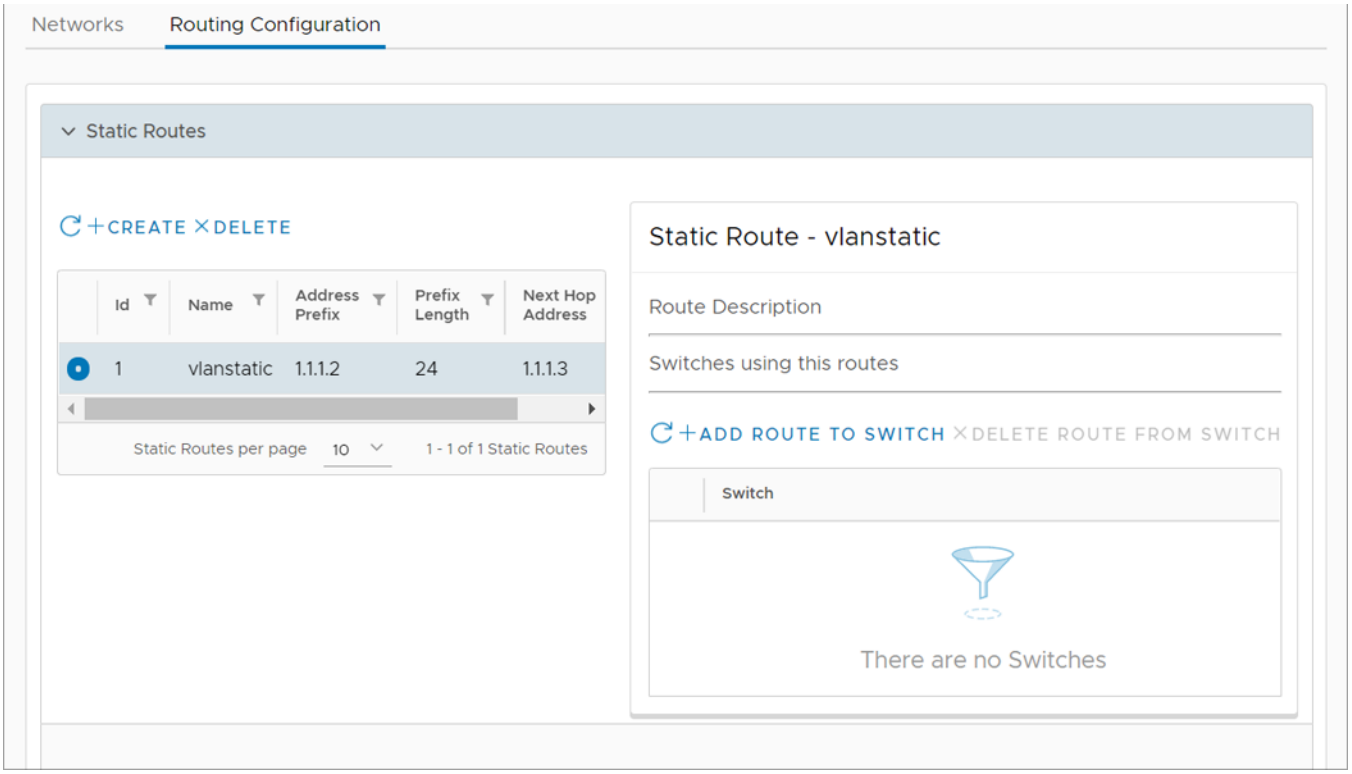

3. Select the switch to map to this route, and click **Add**.

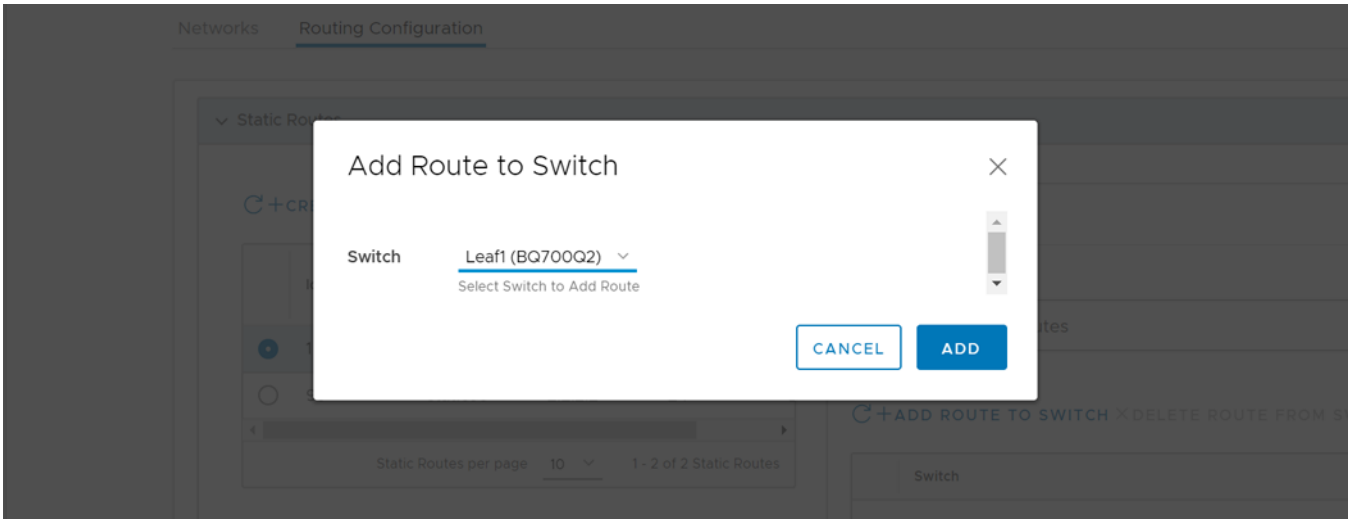

4. The system displays the route added success message.

Static route details

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 static route

1. Select the static route to delete, and click **Delete**.

2. Click **Delete** to confirm.

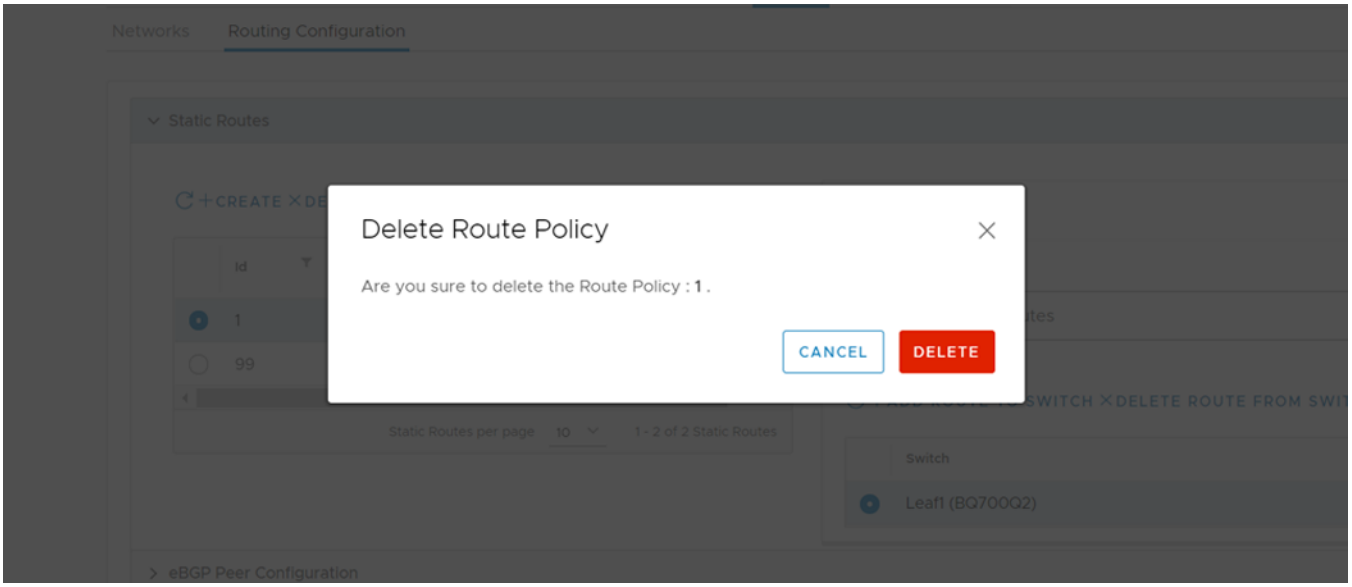

3. The system displays static route deletion is successful.

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.

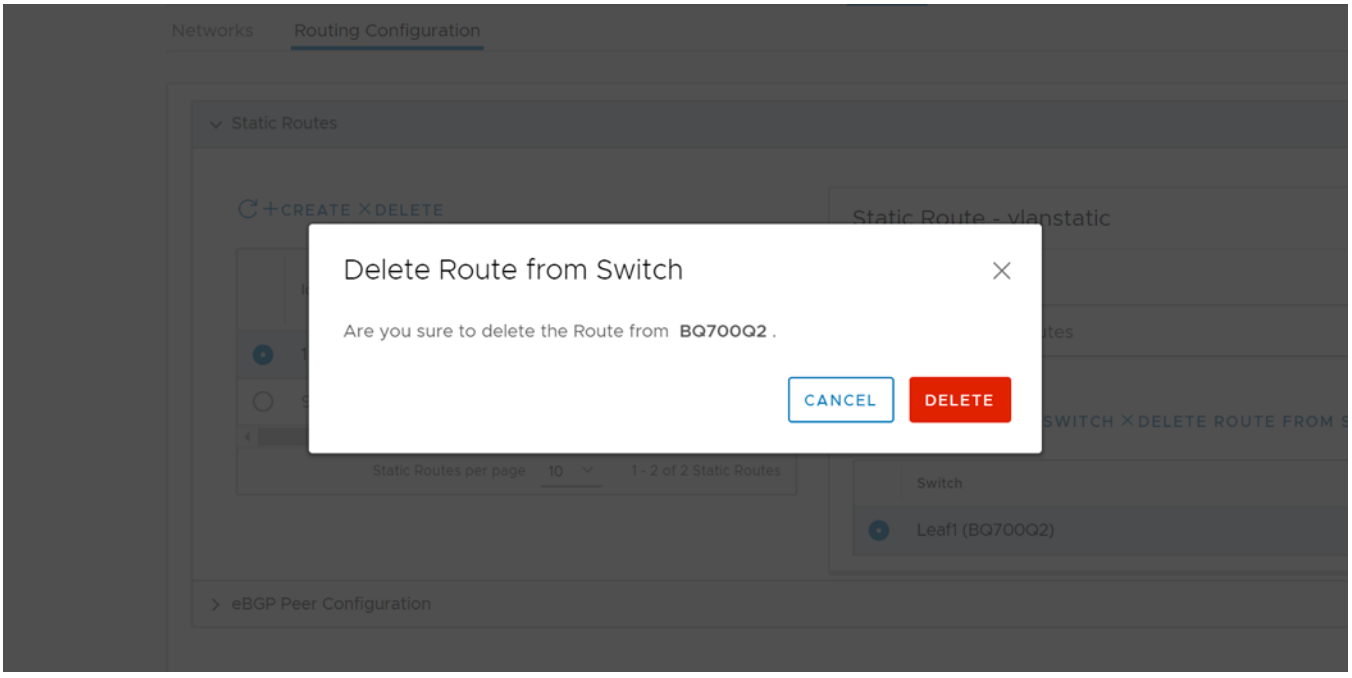

3. The system displays route policy deletion success message.

Configure eBGP peer route

You can configure eBGP peer routes for a network.

Create eBGP route

1. Select **Network** > **Routing Configuration**, and click **eBGP Peer Configuration**.

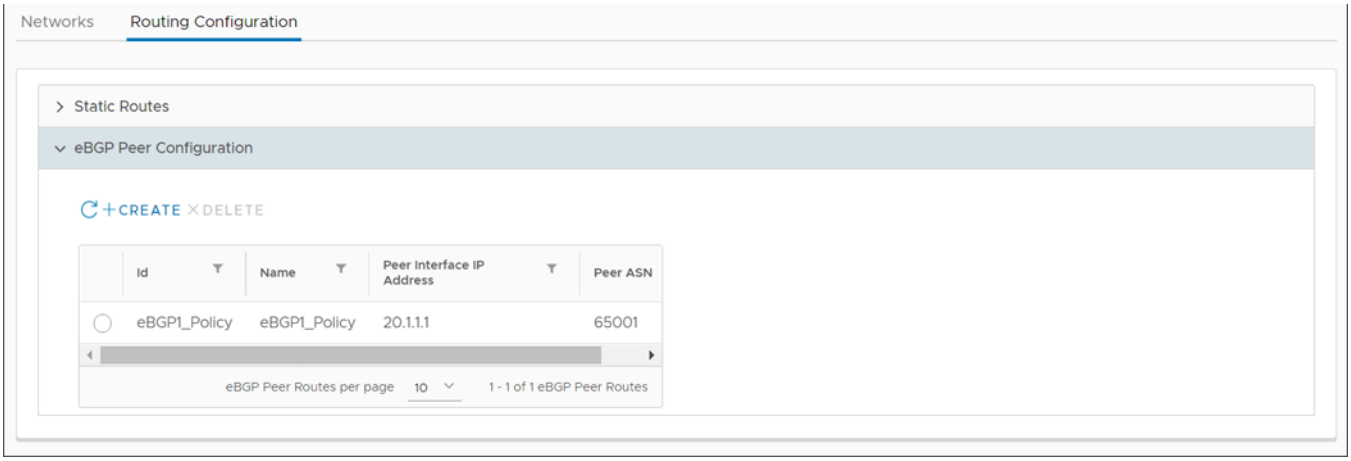

2. Click **Create** to add an eBGP peer route.

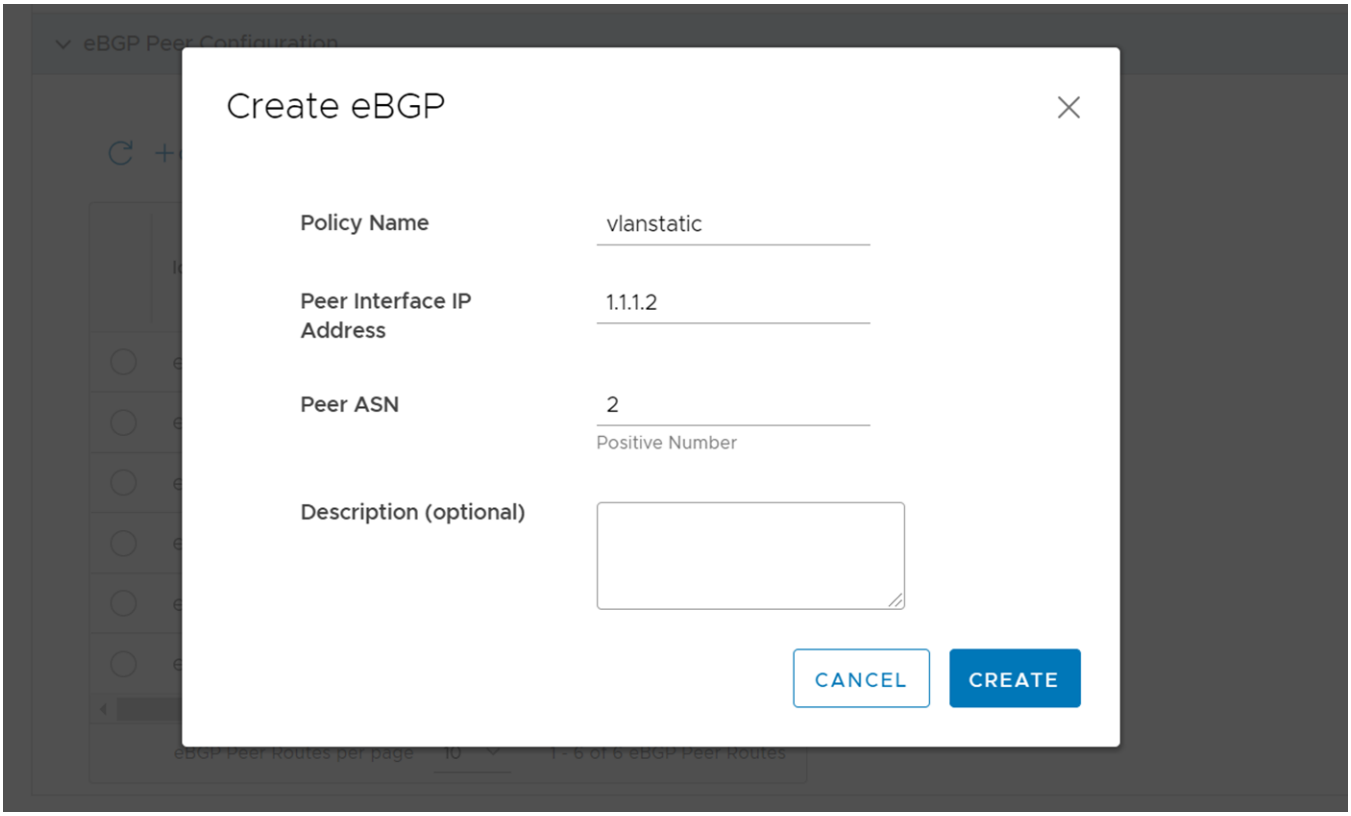

3. Enter the relevant details and click **Create**.
4. The system displays eBGP peer route creation is successful.

Delete eBGP route

1. Select the eBGP route to delete, then click **Delete**.

2. Click **Delete** to confirm.

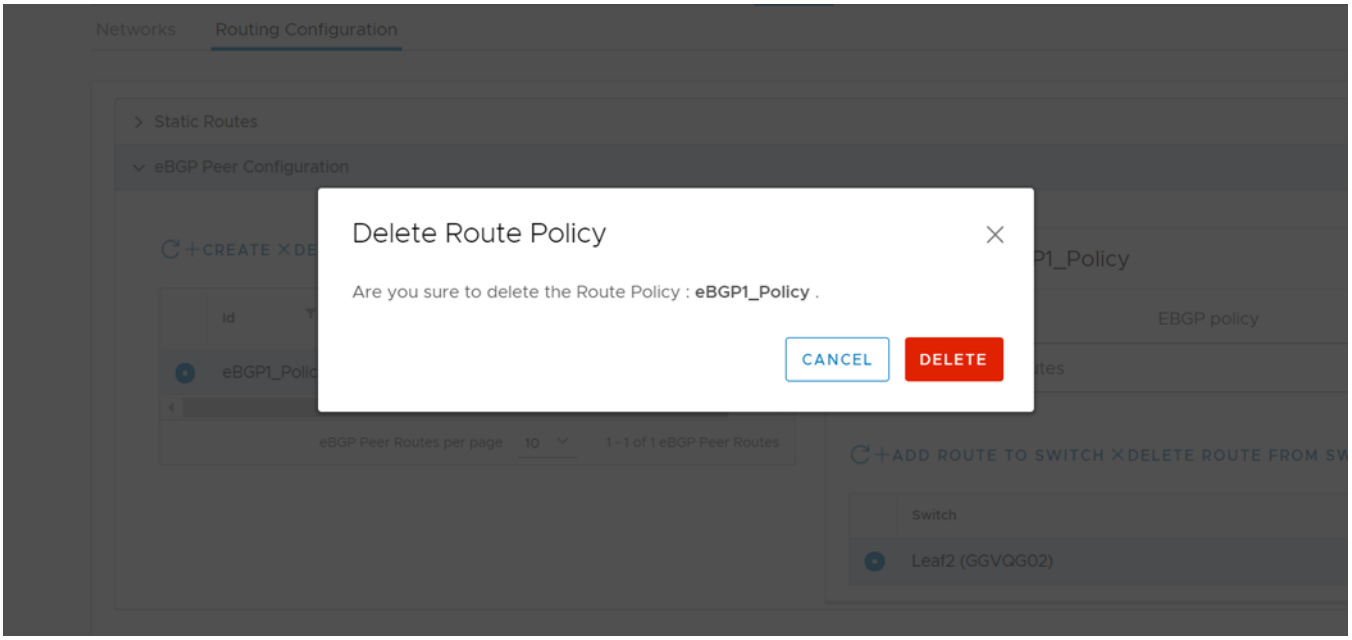

3. The system displays route policy deletion success message.

Add eBGP route to switch

1. Select an eBGP route, then click **Add Route to Switch**.

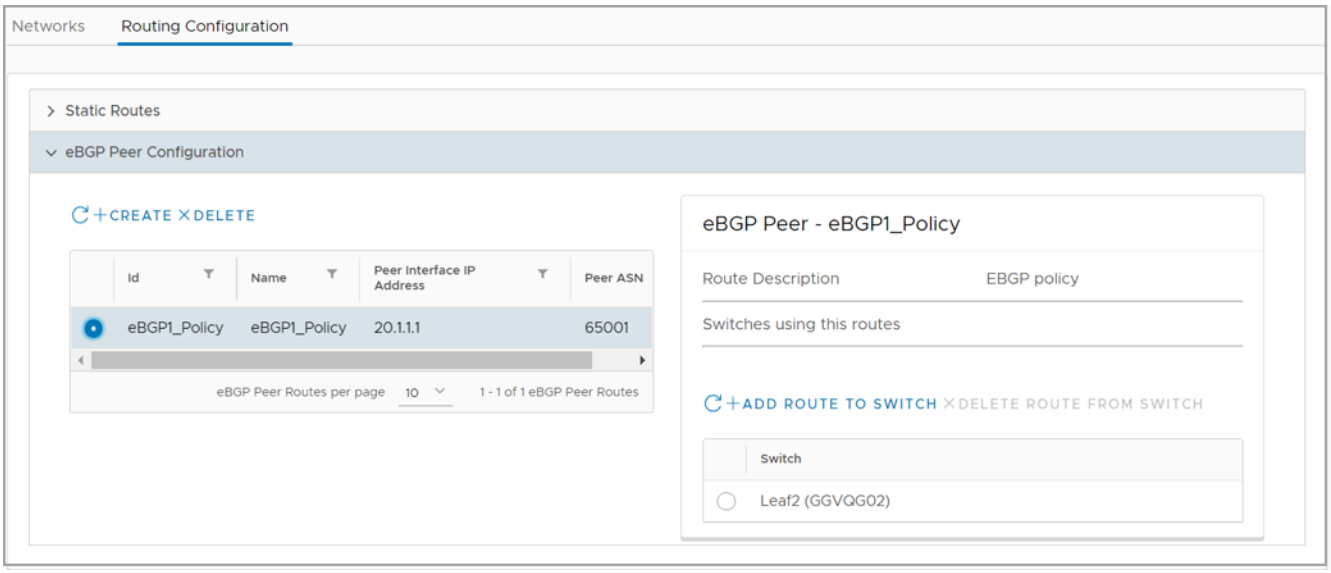

2. Select the switch, then click **Add**.

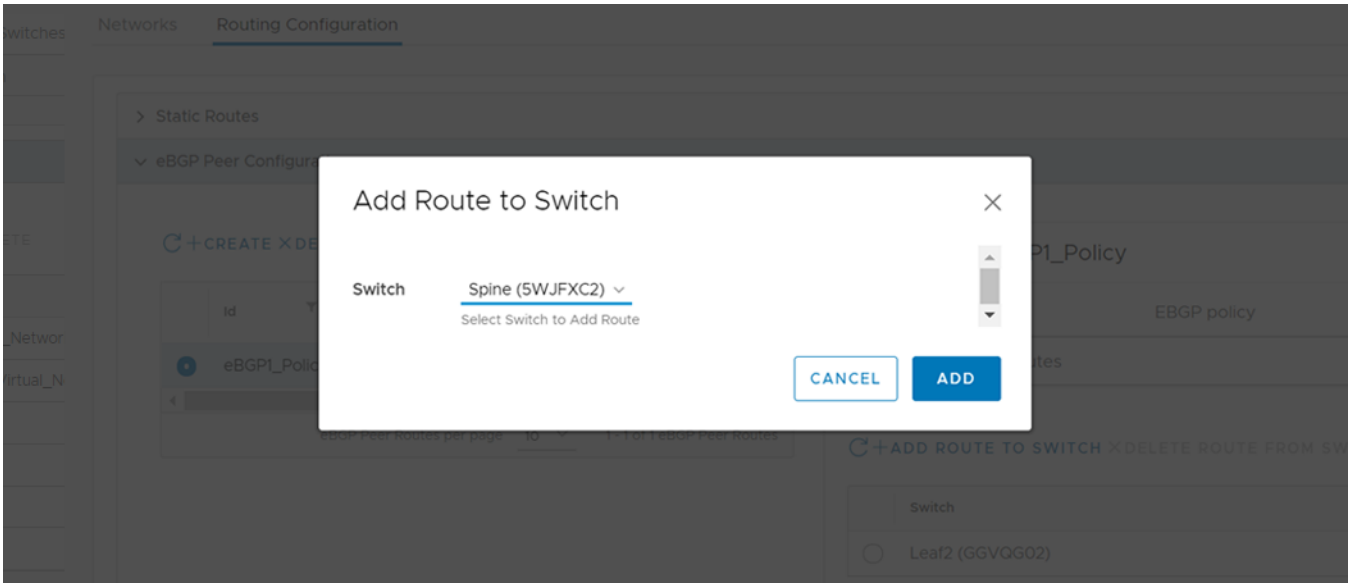

3. The system displays the route to switch addition success message.

View eBGP peer details

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.

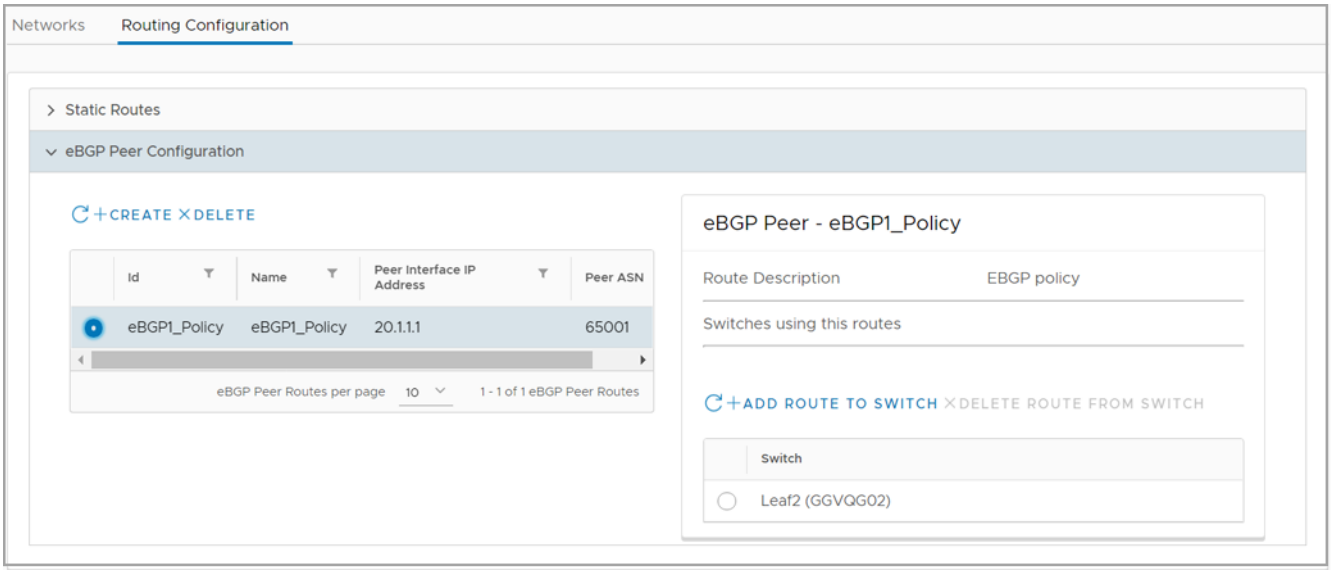

Delete eBGP route from switch

1. Select an eBGP route, then click **Delete Route**.

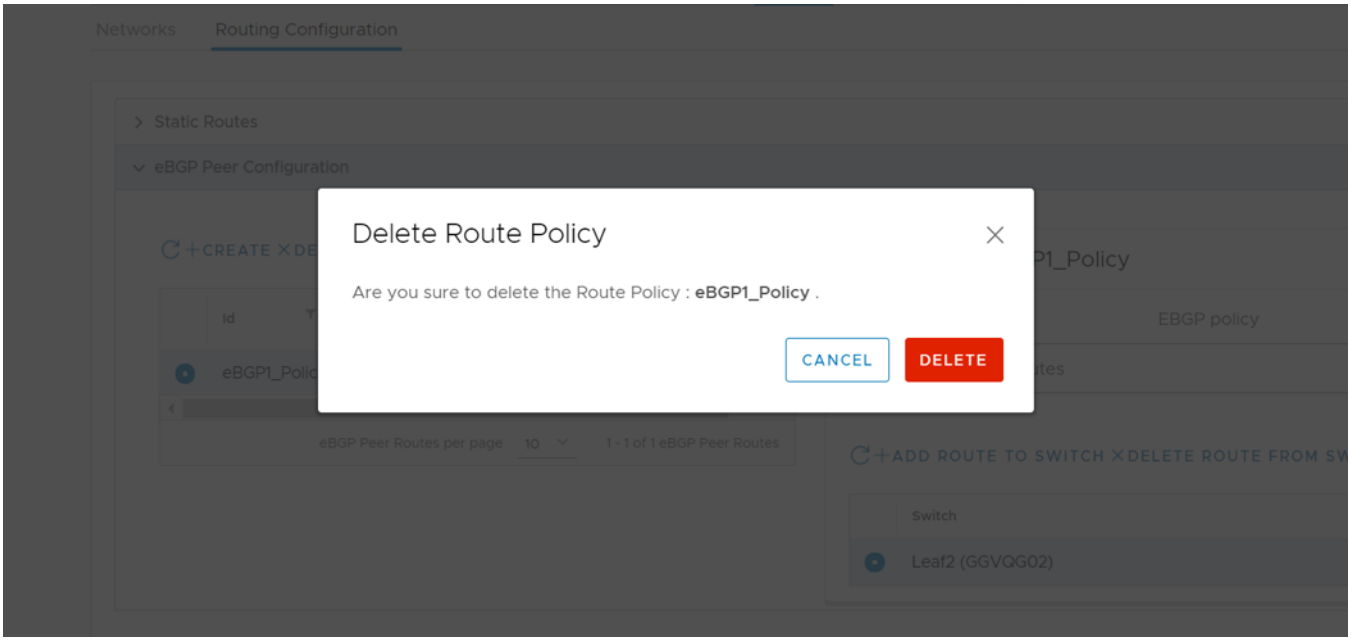

2. Click **Delete** to remove the route from the switch.
3. The system displays route deletion success message.

Configure global settings for SmartFabric

Starting from 2.0 release, OMNI allows you to configure SmartFabric switch services settings using the UI.

You can configure the following services on the SmartFabric switches using OMNI:

- NTP
- DNS
- Syslog
- SNMP

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

Configure NTP server

To configure an NTP server:

1. Select the SmartFabric instance > **Global Settings** > **NTP**. The page displays the list of the NTP servers that are already configured in the OMNI VM.

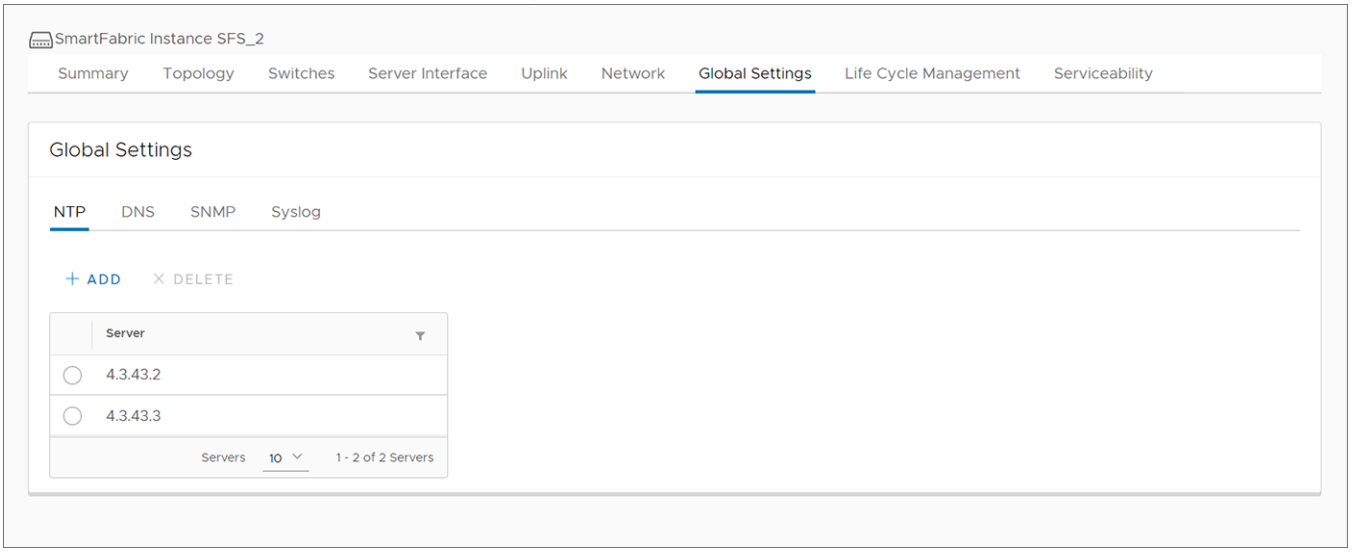

2. Click **Add** to configure an NTP server.
3. Enter the IP address or hostname of the NTP server and click **Add**.

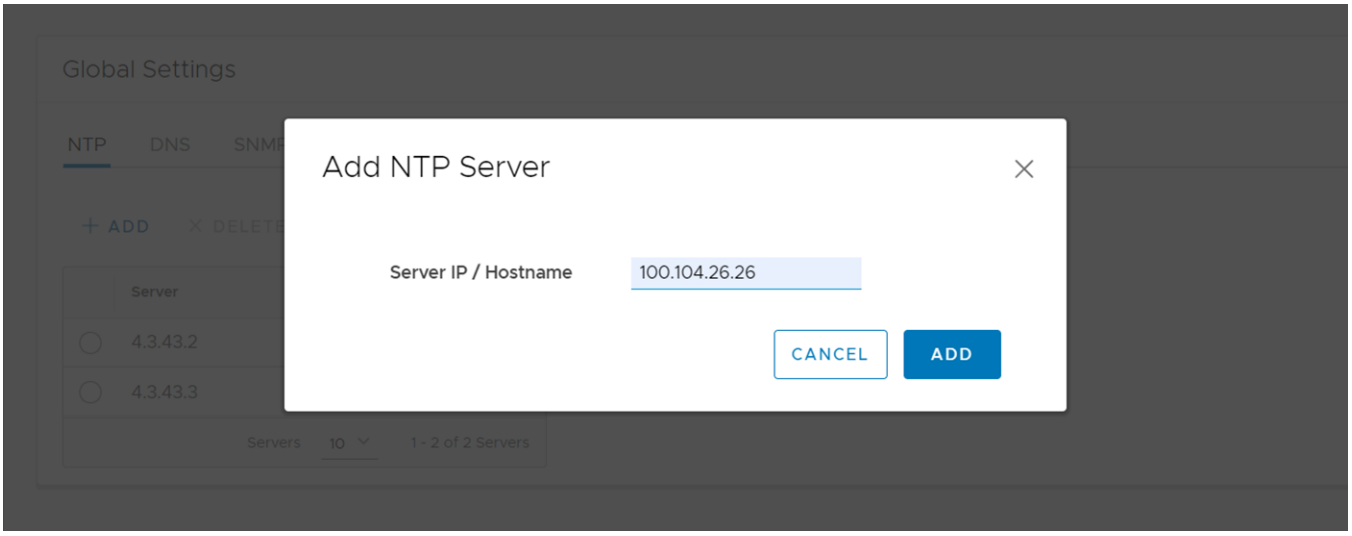

4. The system displays the configuration success message.

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

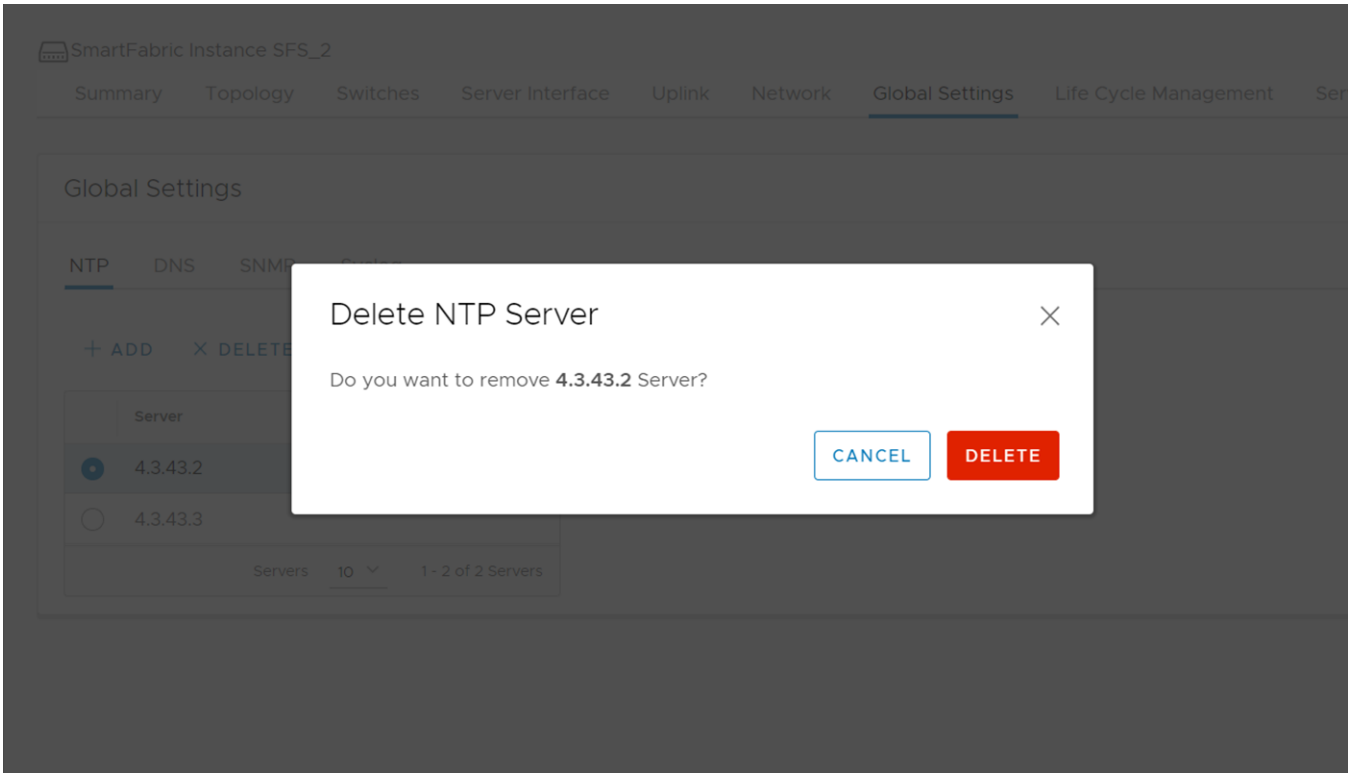

Configure DNS server

To configure one or more DNS servers:

1. Select the SmartFabric instance > **Global Settings** > **DNS**. The page displays the list of the DNS servers that are already configured in the OMNI VM.
2. Click **Add** to configure one or more DNS servers.
3. To configure a single DNS server setting, enter the IP address of the DNS server and click **Add**.

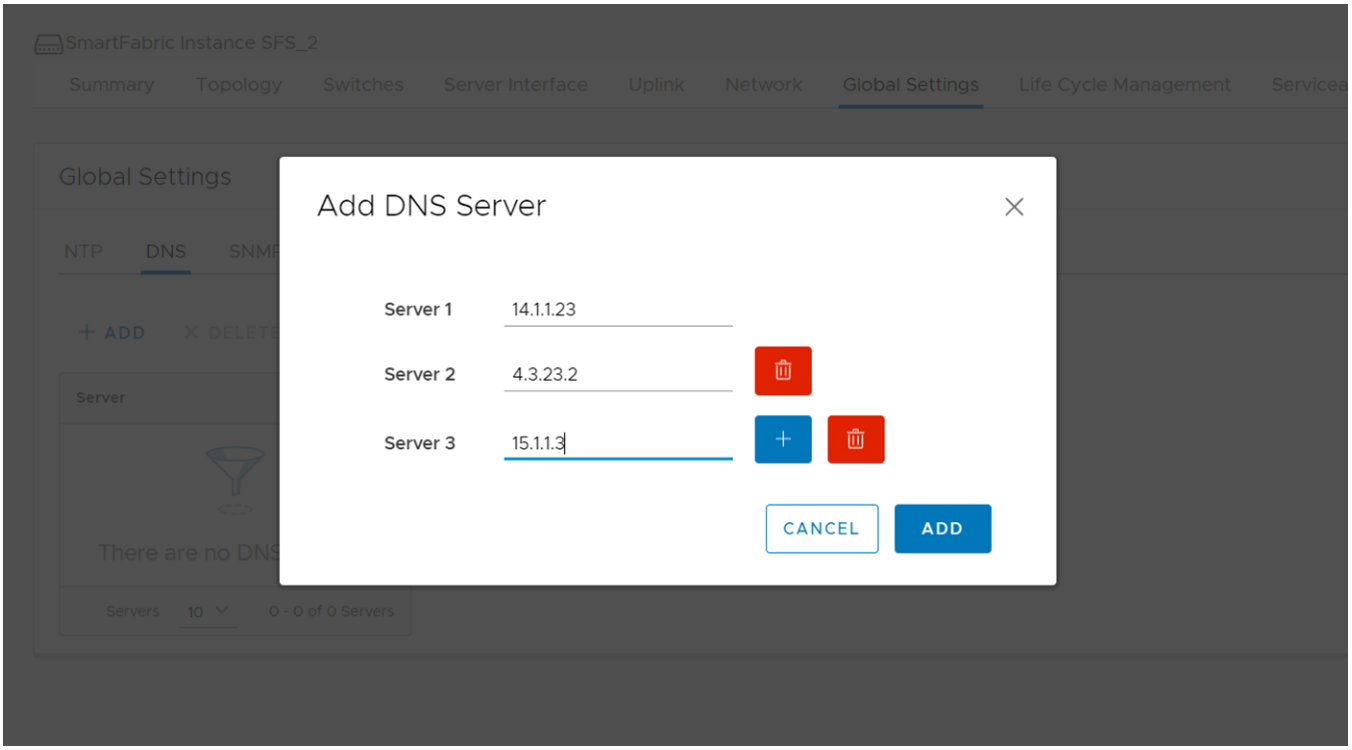

4. To configure multiple servers at a time, enter the IP address of all the servers using + button.

5. The system displays the configuration success message.

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

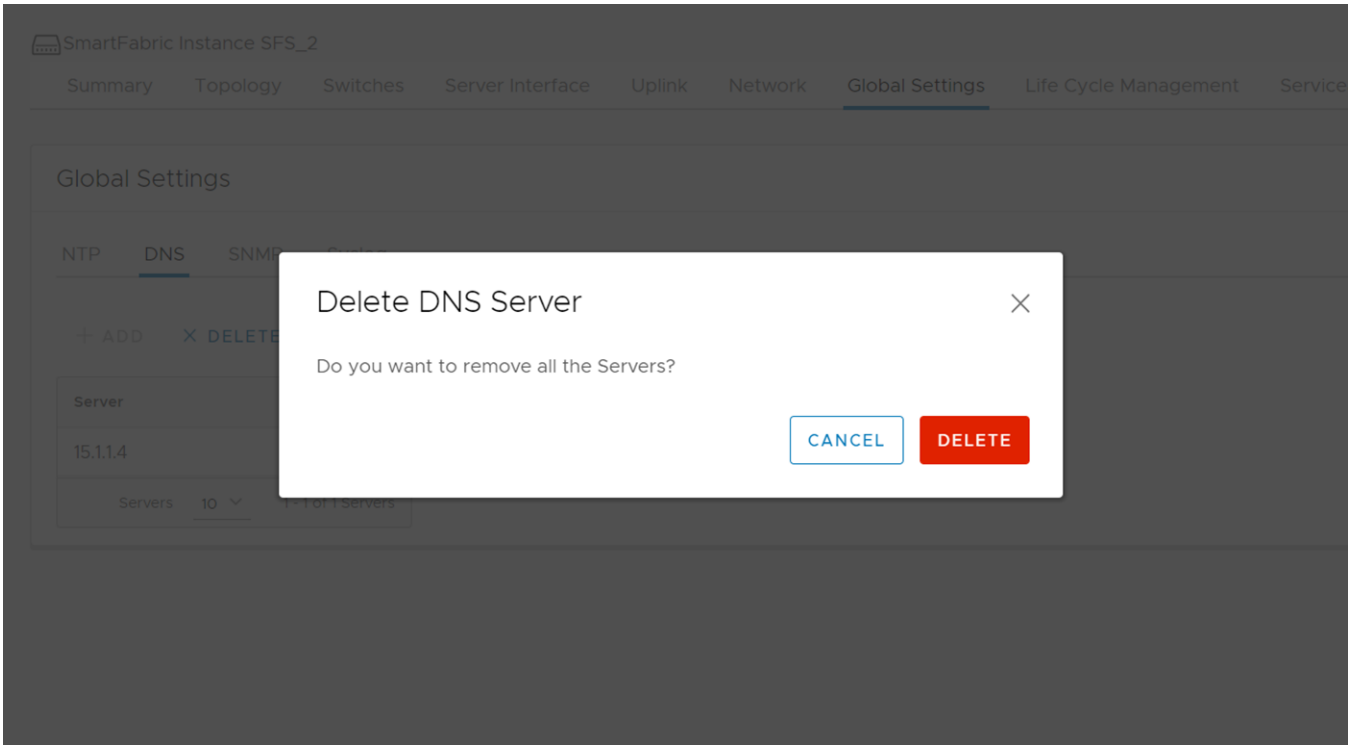

Configure SNMP server

To configure or edit an SNMP server:

1. Select the SmartFabric instance > **Global Settings** > **SNMP**. The page displays the list of the SNMP servers that are already configured in the OMNI VM.
2. Click **Add** to configure an SNMP server.

3. Enter the IP address of the SNMP server, community, SNMP version, and click **Add**.

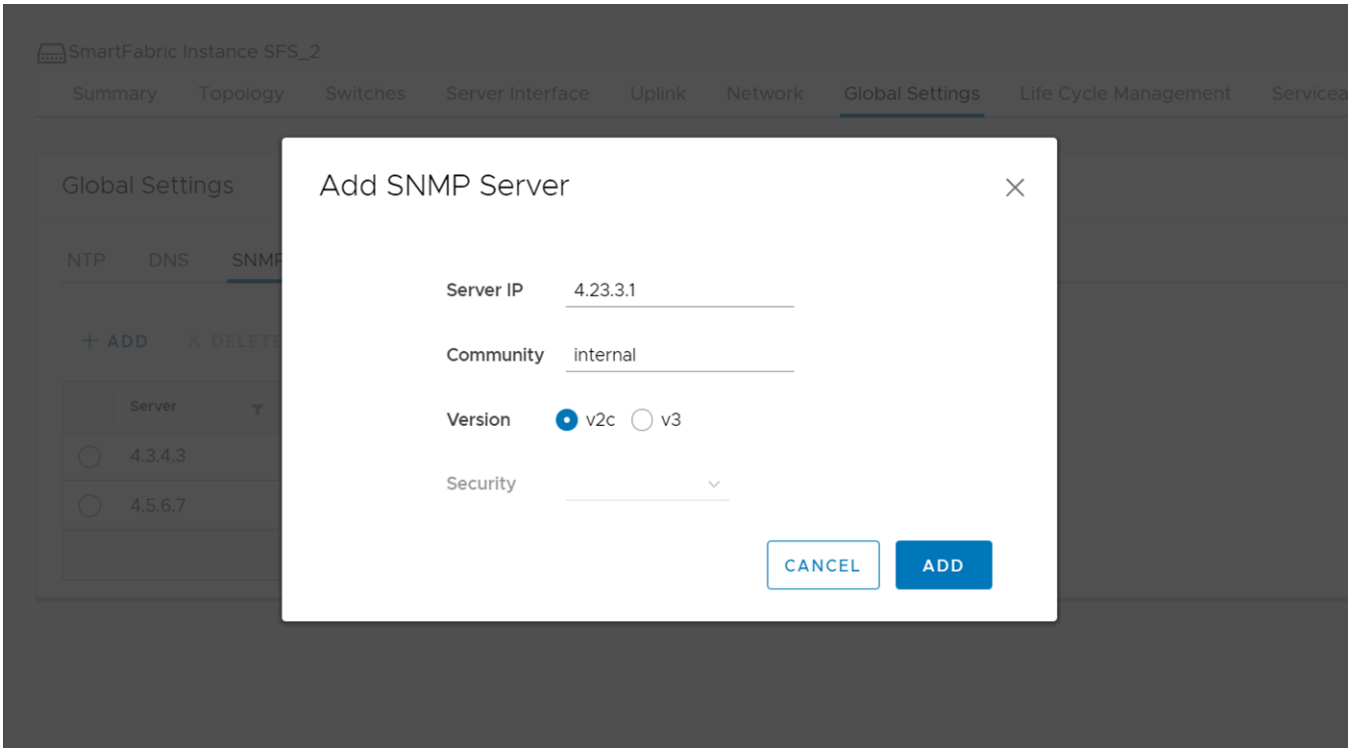

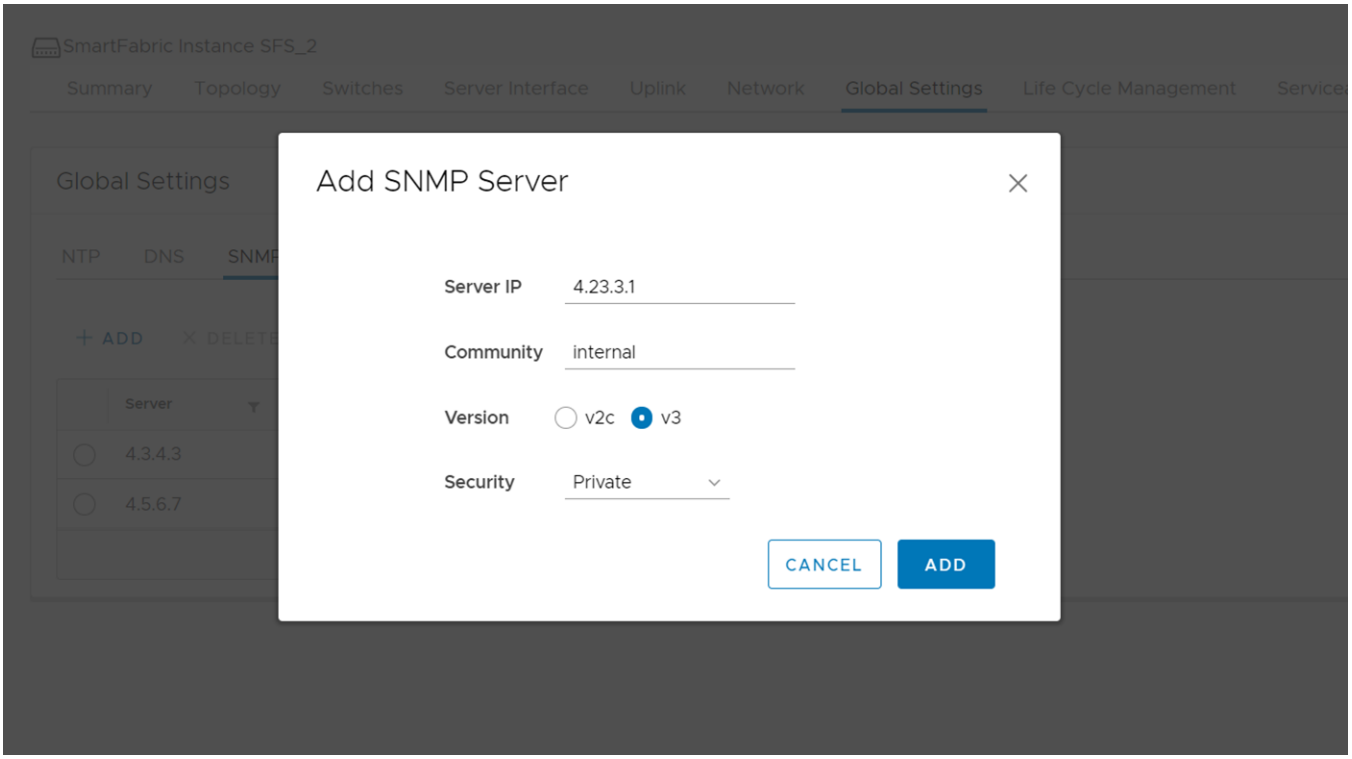

4. The system displays the configuration success message.

To delete the configured servers, click **Delete**.

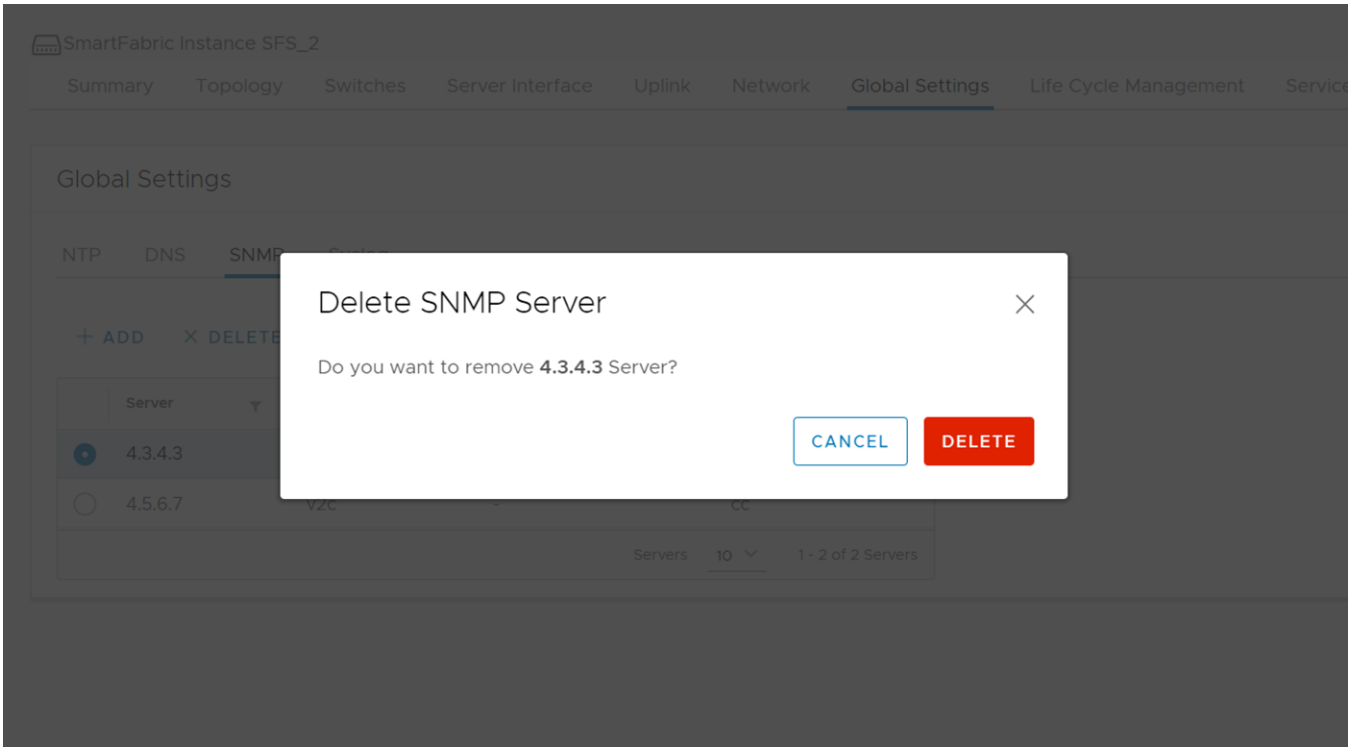

Configure syslog server

To configure and edit a syslog server:

1. Select the SmartFabric instance > **Global Settings** > **Syslog**. The page displays the list of the syslog servers that are already configured in the OMNI VM.
2. Click **Add** to configure syslog server.
3. Enter the IP address of the syslog server and log level, and click **Add**.

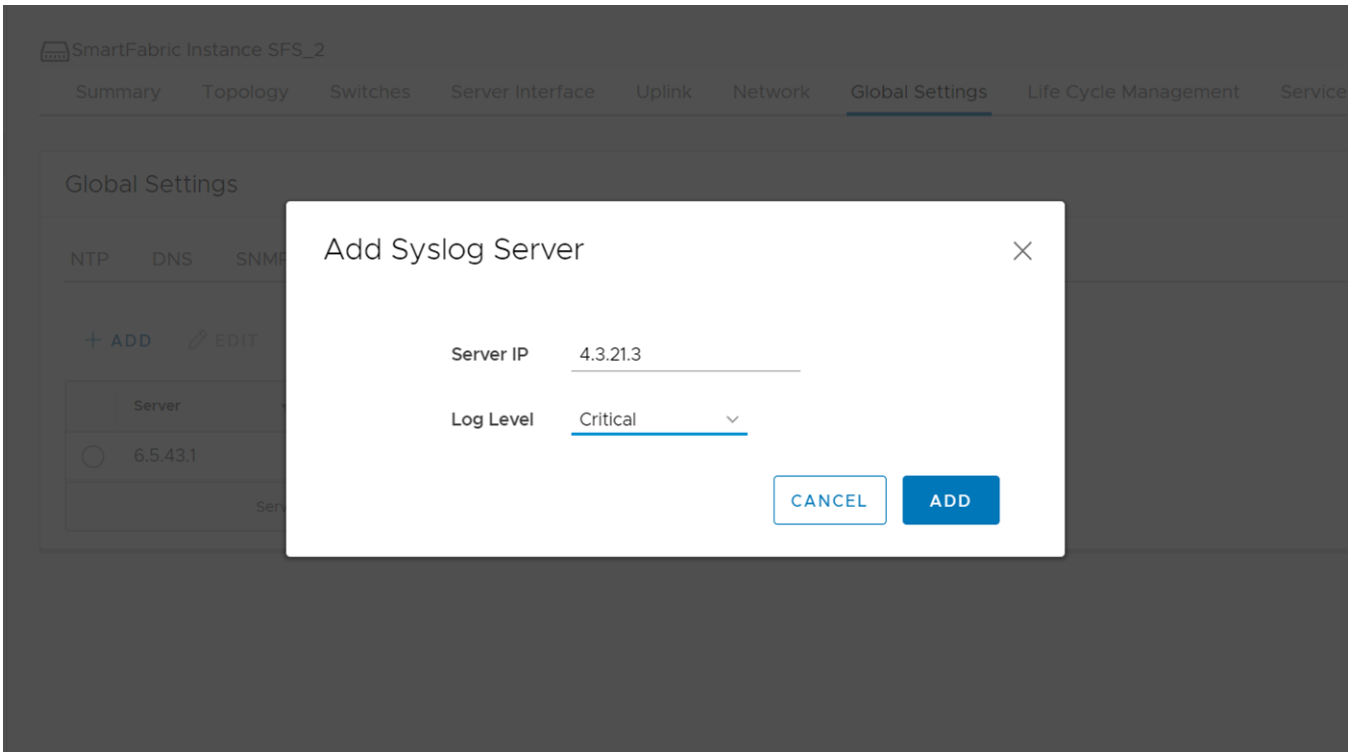

4. The system displays the configuration success message.

Edit syslog server

You can edit the log level for the syslog server. To edit:

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

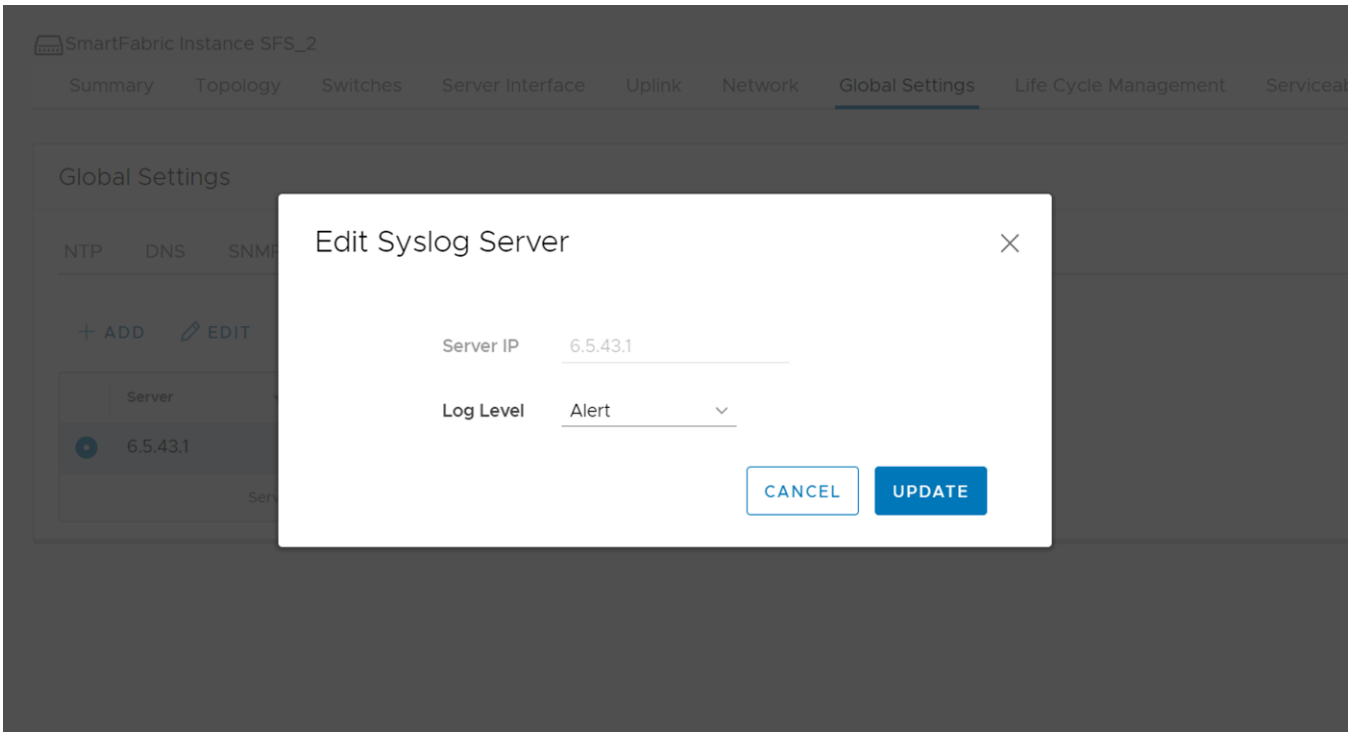

To delete the configured servers, click **Delete**.

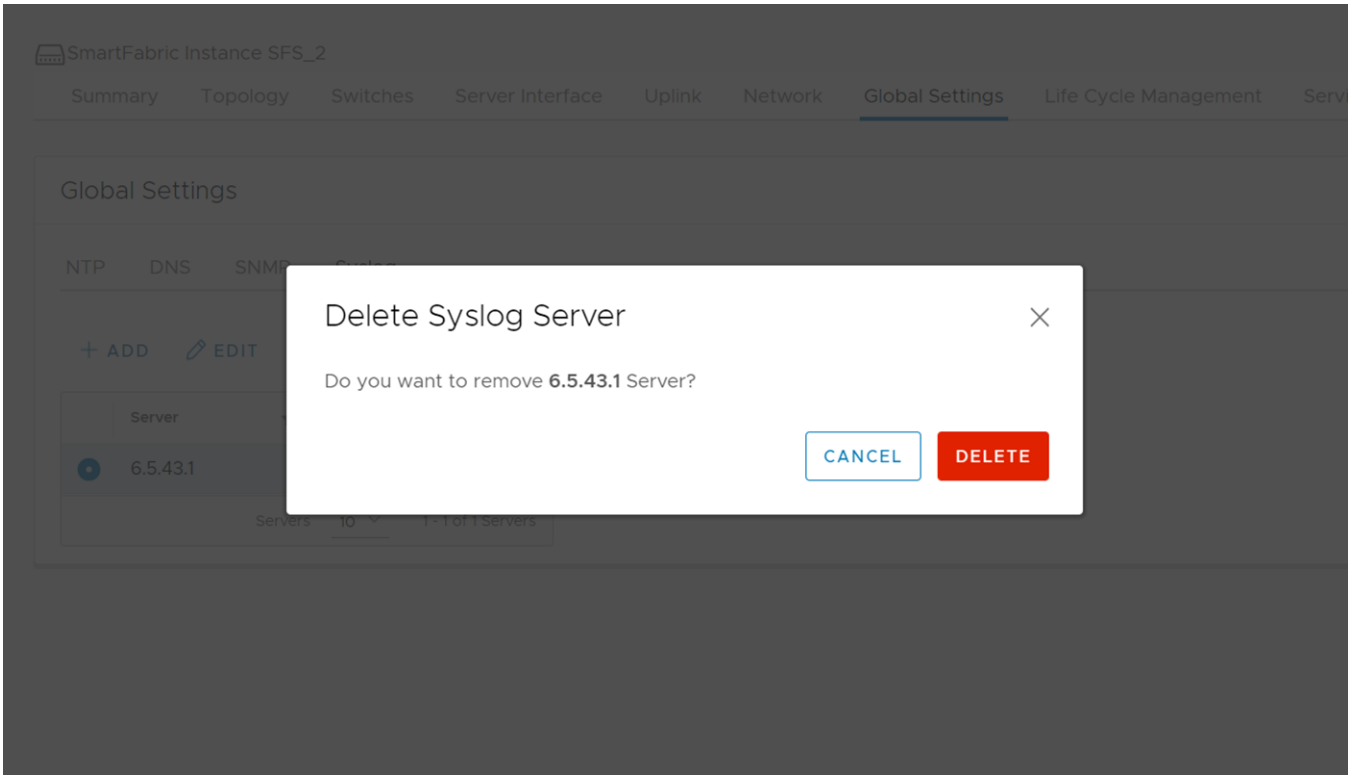

View fabric events and compliance status

Starting from 2.0 release, OMNI displays the list of fabric events and compliance checks for each SmartFabric instance.

Download the events and compliance errors

NOTE: This option is available only when OMNI is accessed as a stand-alone application.

You can download all the events and the compliance errors that are listed for each SmartFabric instance from stand-alone OMNI UI.

Select the SmartFabric instance > **Serviceability** and click **Download**. The downloaded `zip` file contains the fabric events and compliance errors in CSV format.

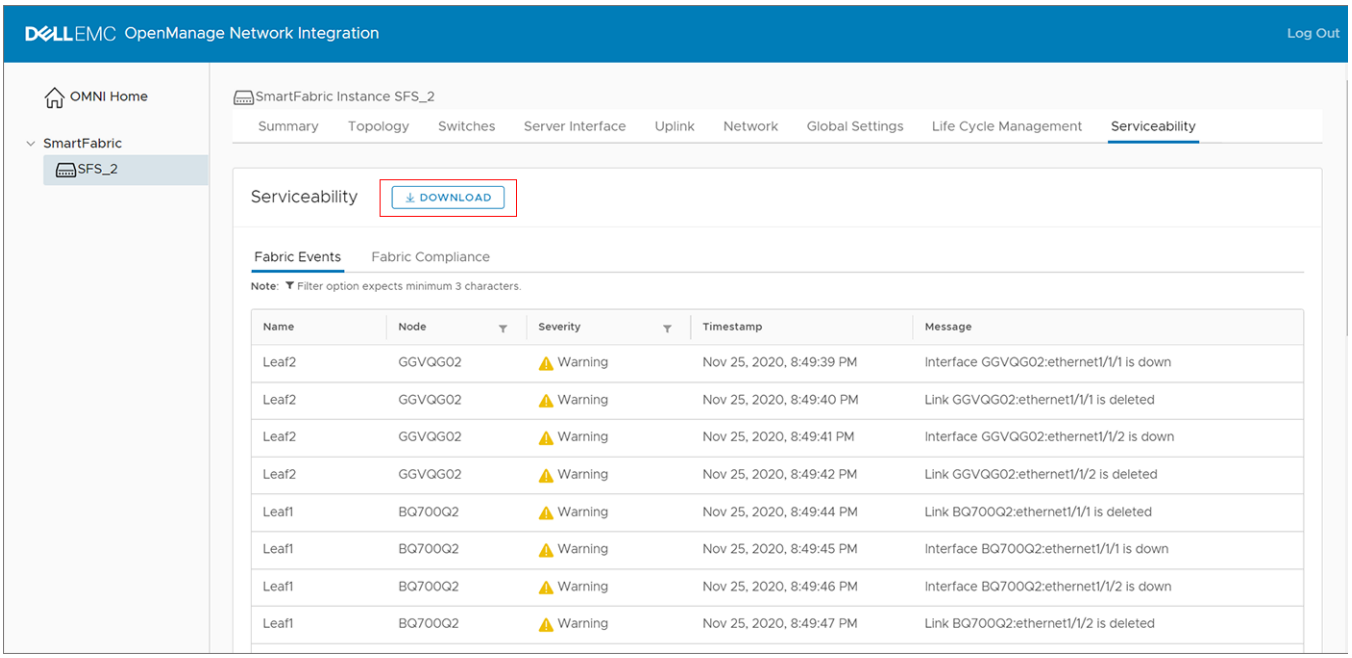

NOTE: **Download** option is not available when OMNI plug-in is launched from vCenter. Hence, you cannot download the fabric events and compliance CSV files from OMNI plug-in page.

View fabric events

OMNI UI lists the events that are generated for each SmartFabric instance.

This feature is supported from SmartFabric OS10.5.0.7 version and both on L2 and L3 personality.

To view the latest events, select the SmartFabric instance > **Serviceability** > **Fabric Events**. The table lists the latest events with detailed information including switch name, service tag of the switch, severity, time, and the event message.

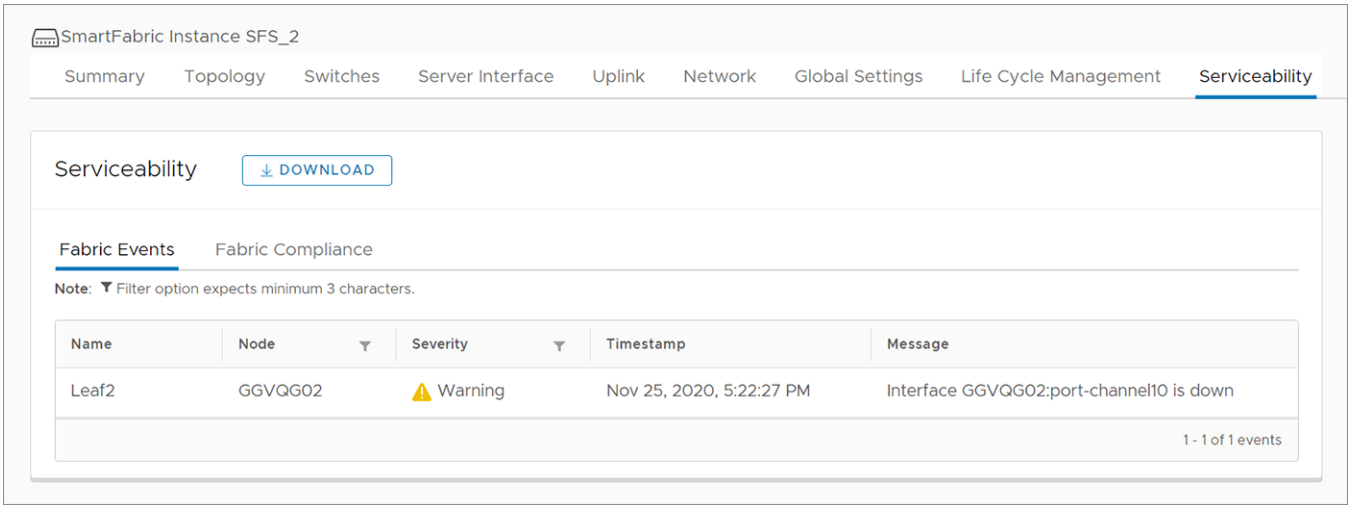

View fabric compliance status

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. OMNI retrieves the fabric compliance status for the SFS instance and displays the noncompliance events with details. OMNI also recommends the actions to eliminate the compliance violations or misconfigurations.

This feature is supported from OS10.5.2.2 version or later, and applicable for SFS L3 leaf and spine personality.

To view the fabric compliance errors:

1. Select the SmartFabric instance > **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.

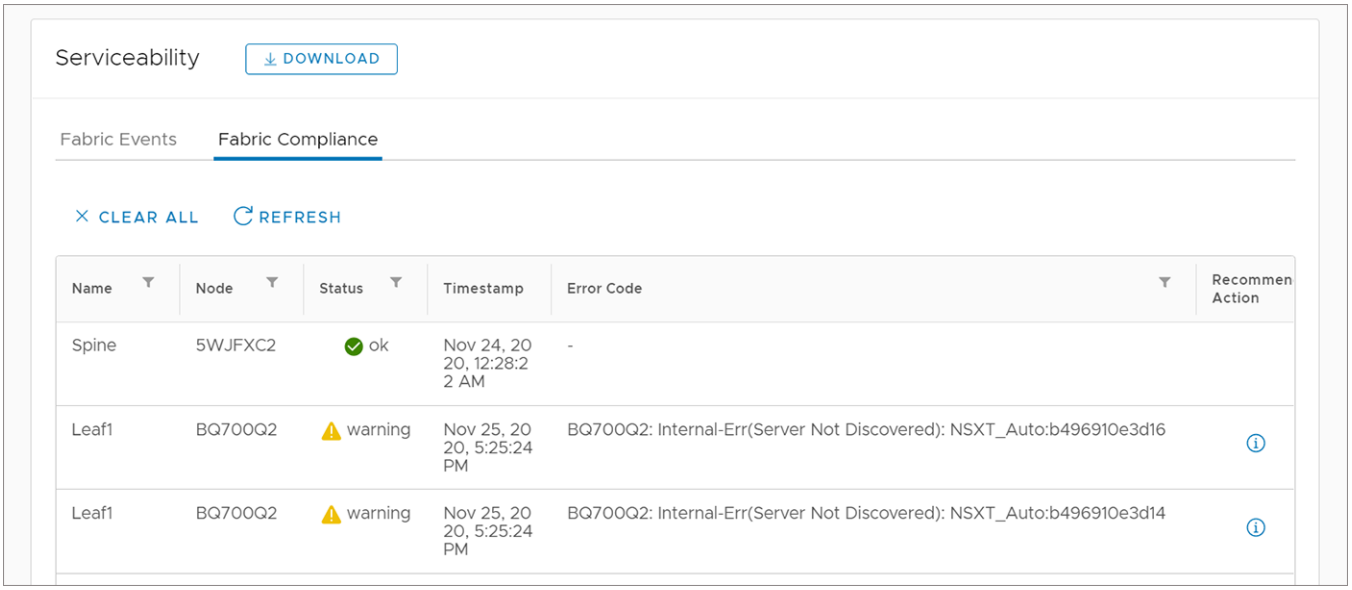

2. Click the information icon to view the recommended action for each compliance error.

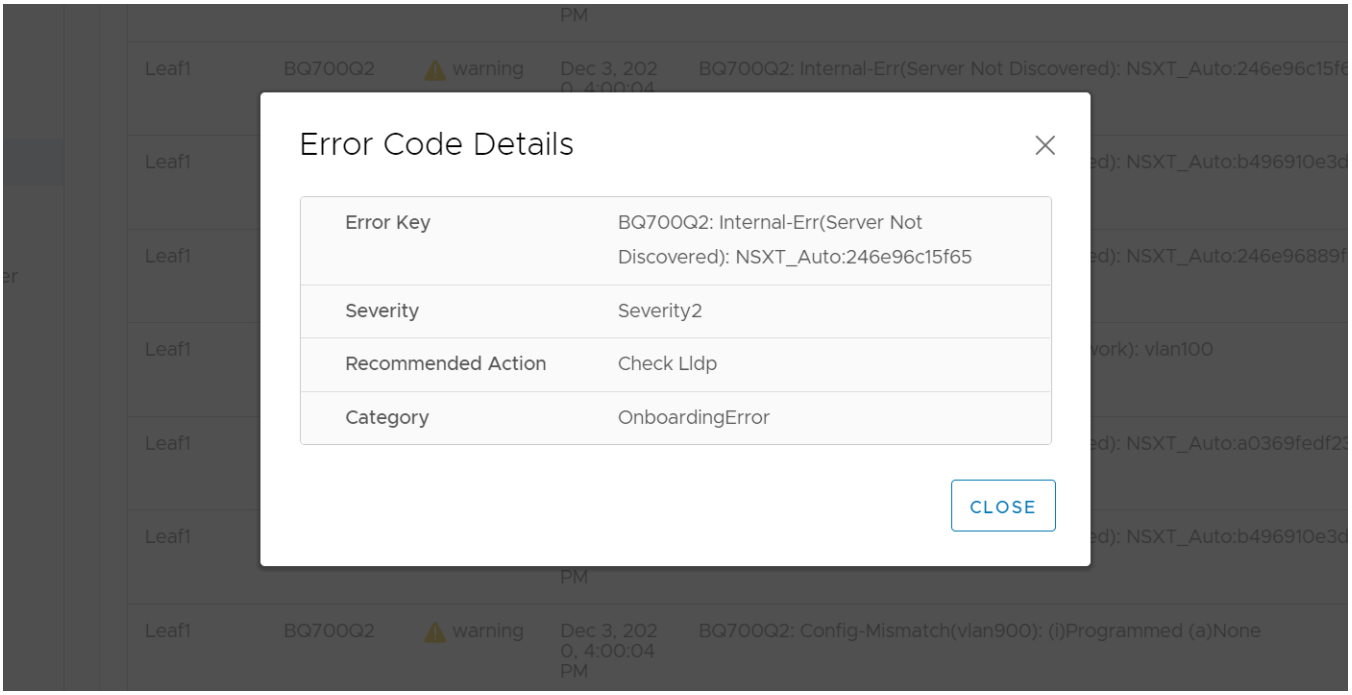

3. (Optional) Click **Clear All** to clear all the existing compliance errors. After clearing all the errors, OMNI starts to retrieve the fabric compliance status for the SmartFabric instance immediately.

Click **Refresh** to update the data and display the new compliance errors.

You can also view the fabric events and the compliance errors in the SmartFabric instance overview dashboard. Select the SmartFabric instance > **Summary** > **Overview** to view the overview of events and errors. The fabric compliance errors are grouped under infrastructure, cluster, server onboarding, and uplink categories.

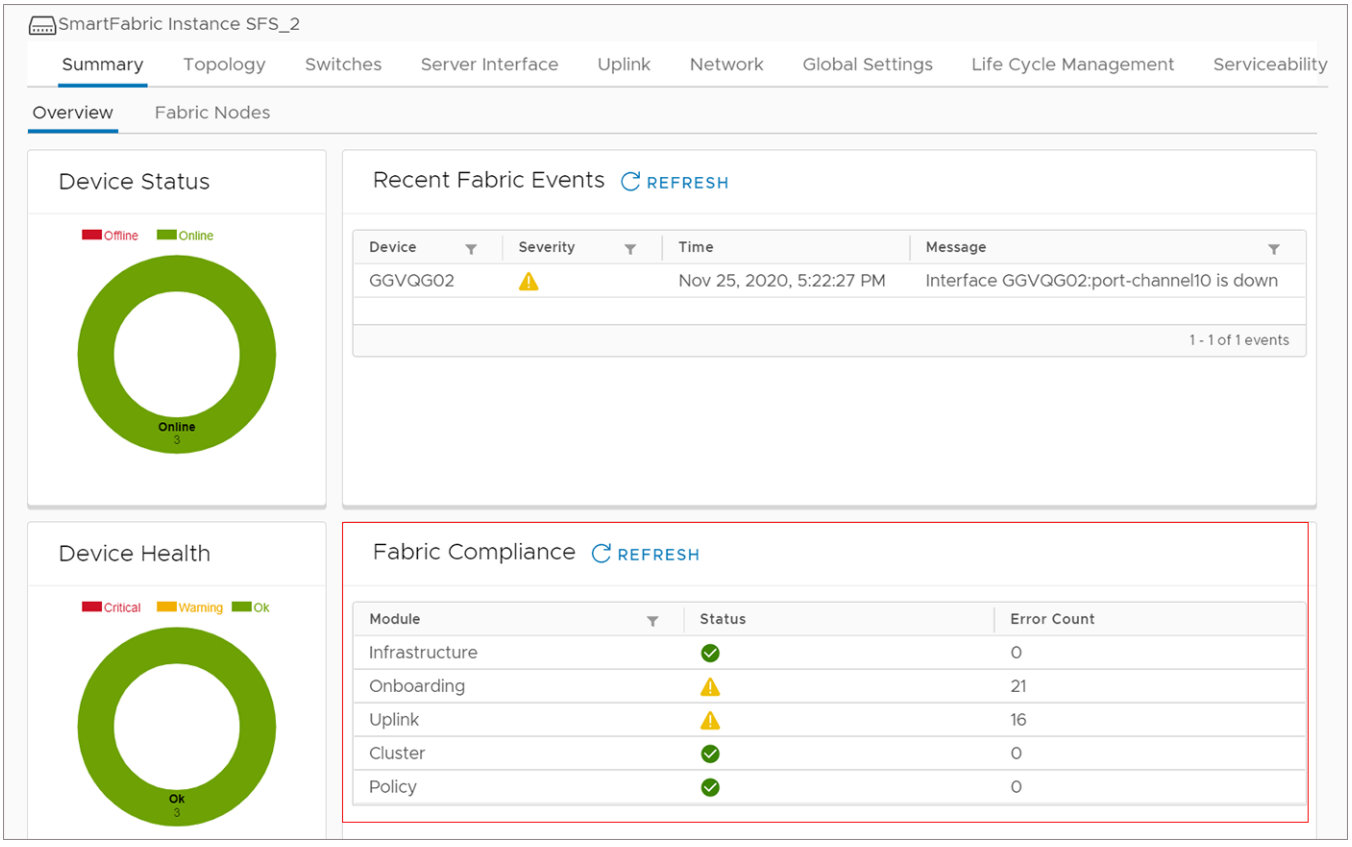

10

Lifecycle management

This chapter explains common lifecycle operations of upgrading the SmartFabric OS10 and OMNI appliance, replacing a switch, and backup and restoring the SmartFabric.

NOTE: The Lifecycle management features are not supported on OME-Modular instances. For more information, see OMNI feature support matrix.

Using **Life Cycle Management** menu, you can:

- Change SmartFabric password.
- Upgrade SmartFabric OS10 image.
- Replace a switch in a network fabric.
- Fabric backup and restore.

Change SmartFabric password

To change the SmartFabric password:

1. Select the SmartFabric instance > **Life Cycle Management** > **SmartFabric Password Change**.

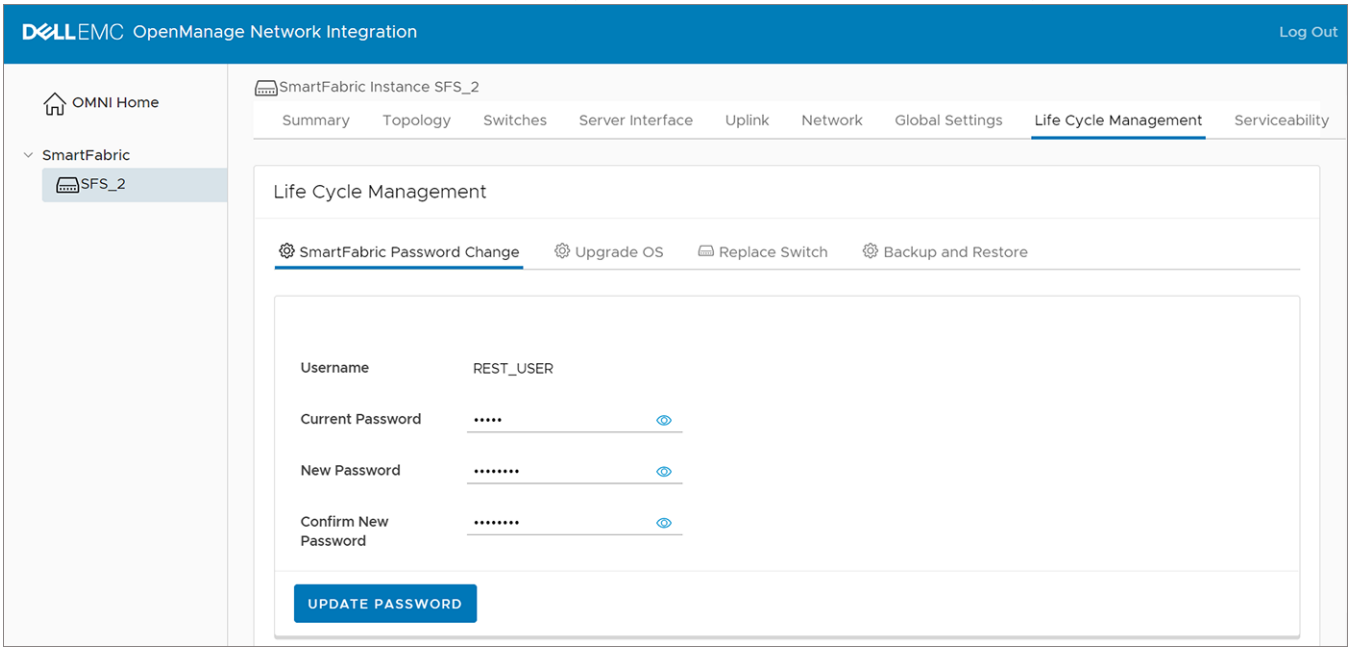

2. Enter the current password for the REST_USER, the new password, confirm the new password, and click **Update Password**.
3. The system displays password update success message.

Upgrade SmartFabric OS in switch

You can upgrade SmartFabric OS from OMNI VM.

From OMNI, you can upload an OS10 image to upgrade the switches in SmartFabric.

NOTE: This instruction is applicable for SmartFabric instance. To upgrade OS10 on MX switches, use OME-M console. For more information, see PowerEdge MX documents.

You can upgrade OS using the following steps:

- Upload the latest image in the OMNI VM.
- Upgrade fabric using the uploaded image.
- (Optional) Delete the image from the OMNI VM.

Upload image

Upload an OS10 image to the OMNI VM:

1. Select the SmartFabric instance > **Life Cycle Management** > **Upgrade OS**.

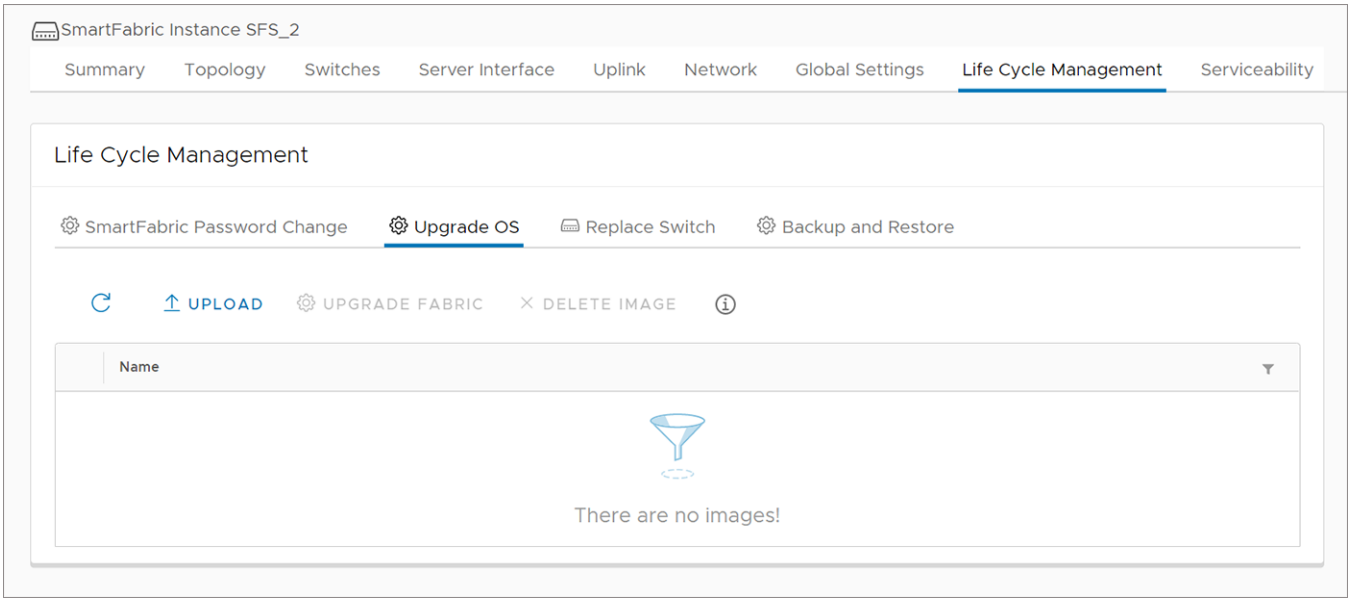

2. Click **Upload** to upload the .bin file.

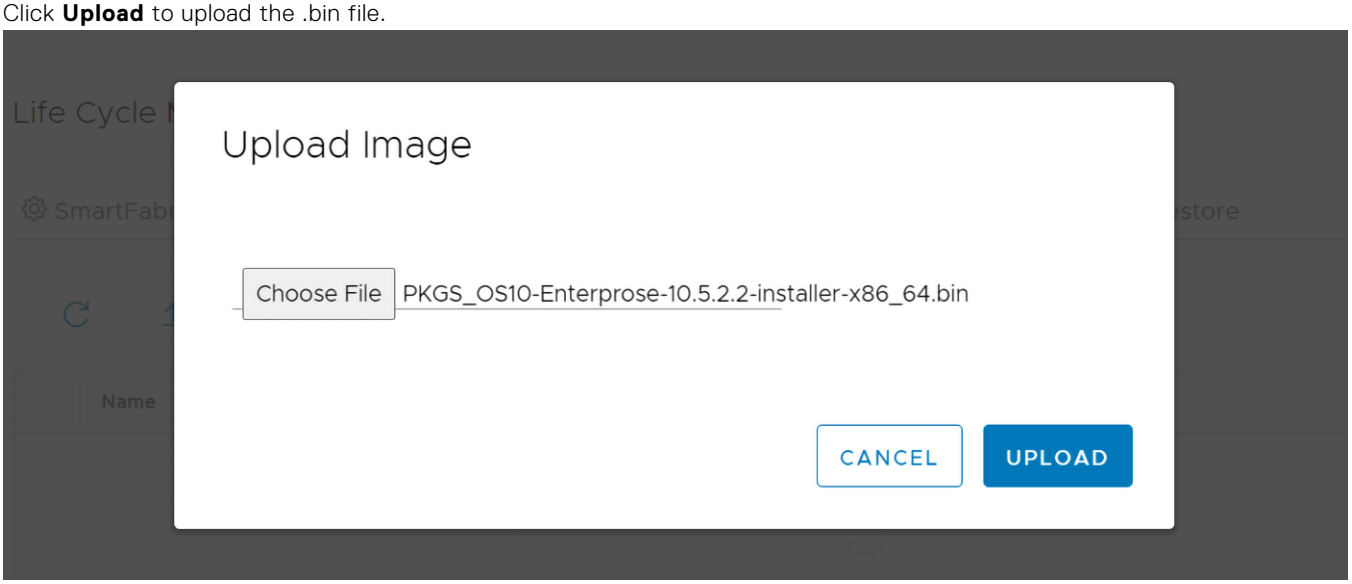

Upgrade fabric

Click the informational icon to see the current SmartFabric OS version.

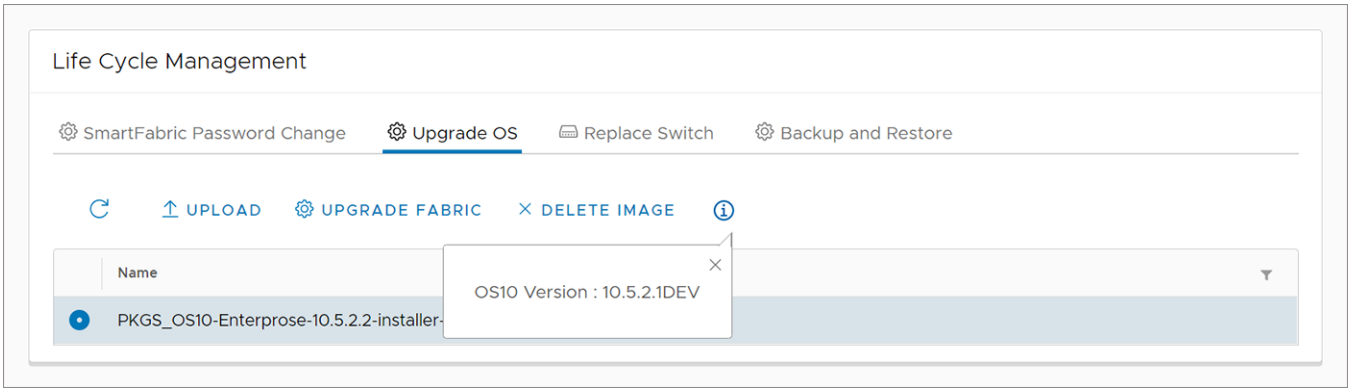

Upgrade the switches in a fabric with an OS10 image:

1. Select the .bin image, and click **Upgrade Fabric**.

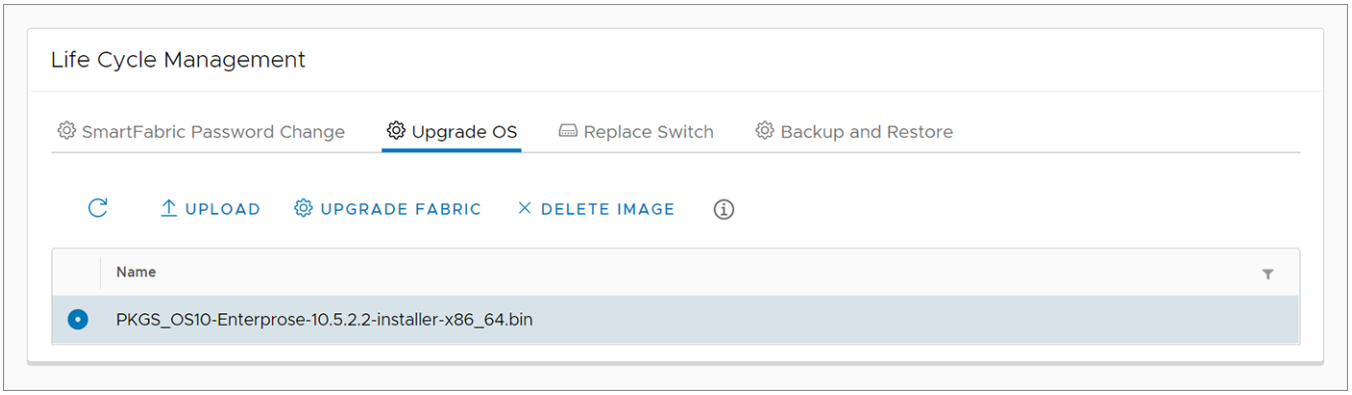

NOTE: Upgrade Fabric option upgrades all the switches in a network fabric. You cannot stop the upgrade after it is triggered.

2. Click **Upgrade** to confirm.

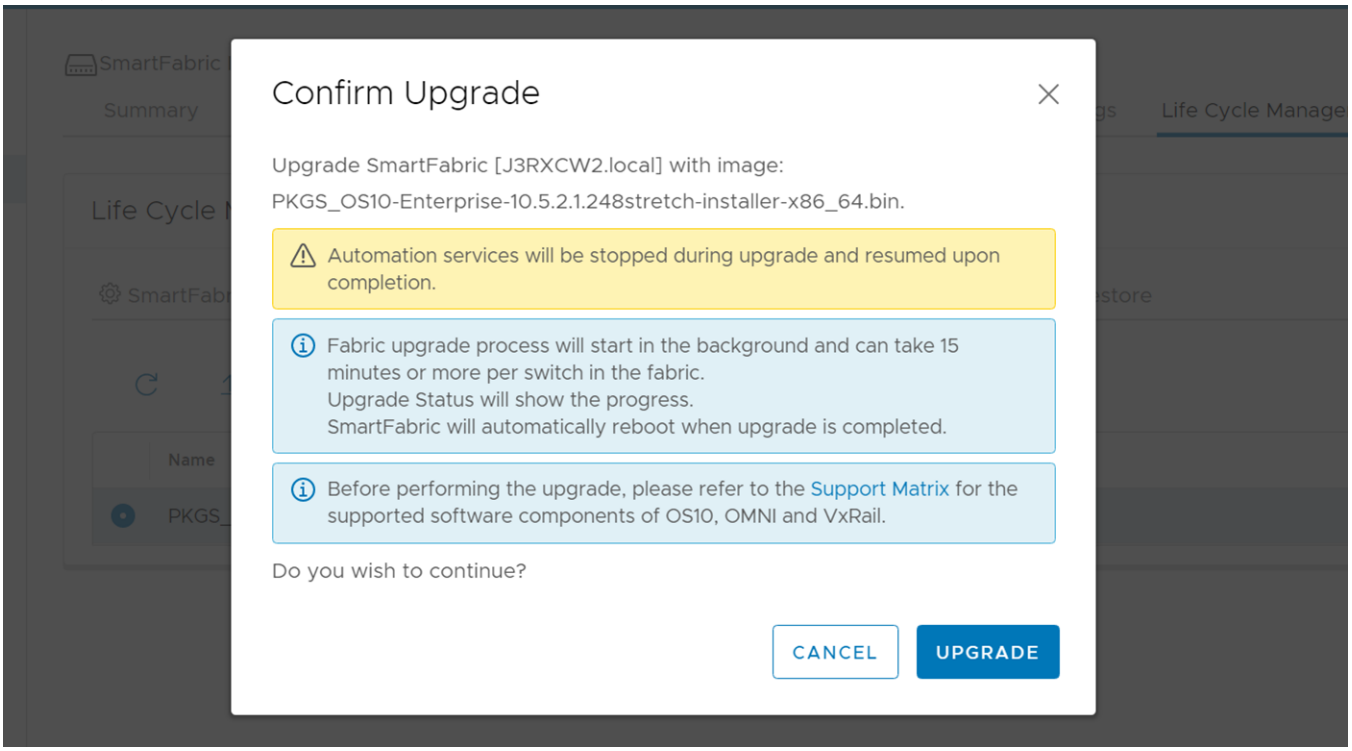

3. The system displays fabric upgrade success message.

   SmartFabric automatically reboots when the upgrade is complete.

Delete image

Delete the OS10 image uploaded in the OMNI VM:

1. Select the .bin image to delete and click **Delete Image**.

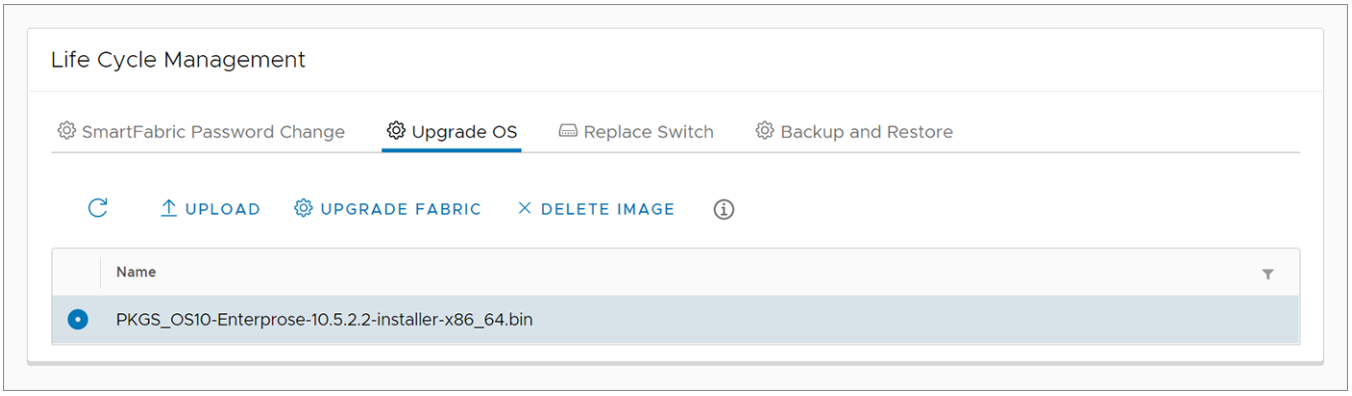

2. Click **Delete** to confirm.

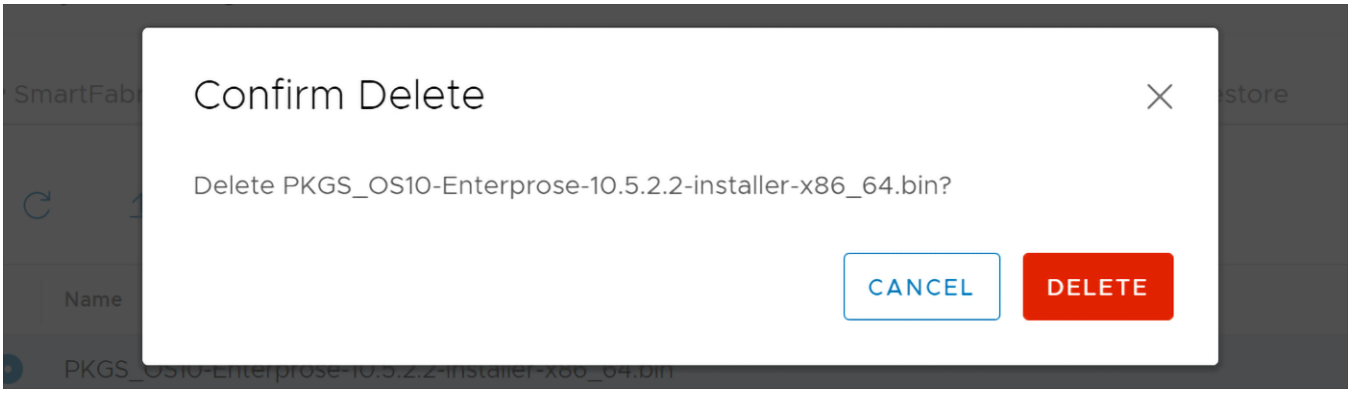

3. The system displays delete image is success.

Replace switch in a fabric

You can replace the faulty OS10 switch in a SmartFabric.

NOTE: This instruction is applicable for SmartFabric instance only. To replace a switch in MX, follow the instructions that are provided in the MX documents. For more information, see PowerEdge MX documents.

To replace a switch:

1. Identify the OS10 switch to be replaced and label each of the cables with the port numbers before disconnecting the cables.
2. Back up the following configurations from the faulty switch to configure the new switch with the same details:
   - Hostname
   - Management IP address
   - DNS and NTP IP addresses if configured
   - Spanning-tree mode

     NOTE: In SmartFabric Services mode, RPVST+ is enabled by default on the uplink interfaces.
   - Other nonfabric commands
3. Ensure that the new switch has the same OS version as the faulty switch. You can check the version using the following command:

```
OS10# show version
```

4. Power off the existing switch to prevent data traffic loss in the cluster.
5. Remove the ICL and uplink connections from the existing switch, and connect to the new switch.

   NOTE: Do not remove connections to VxRail nodes until the new switch is in SmartFabric Services mode.

NOTE: Ensure that the ICL ports are connected to the other leaf switch which is already in SmartFabric Service mode.

6. Enable SmartFabric Services on the new switch and define the ICL ports.
   - For **L2 personality**—Enable SmartFabric Services on the new switch, and define the breakouts, uplinks, interlink ports, plus any other parameters such as management VLAN, LACP, VLAN tagging, and so on.

     For example, if the uplink port is 1/1/4 and the interlink ports are 1/1/29,1/1/30, no VLAN tagging, LACP auto, management VLAN 1 as default.

     ```
     :~$ sfs_enable_vxrail_personality.py -i 1/1/6,1/1/8 -u 1/1/4 -l
     ```

   - For **L3 personality**—Enable SmartFabric Services on the new switch using the `smartfabric l3fabric enable role` command. Example:

     ```
     OS10# smartfabric l3fabric enable role LEAF vlti ethernet 1/1/29-1/1/30
     ```

     For more information about enabling SmartFabric Services, see *Dell EMC SmartFabric OS10 User Guide Release 10.5.0*.

7. The new switch reboots and is placed in SmartFabric Services mode.

   NOTE: During reboot, the configurations are synchronized in the new switch and it takes several minutes.

8. Connect VxRail server ports to the new switch one-by-one to bring up the switch ports and advertise LLDP.
9. Review the command outputs on both switches for same configurations. Use the following commands to validate the configurations:
   - `OS10# show vlan`

     NOTE: The command displays if the switch is a primary or secondary peer.
   - `OS10# show vlt 255`
   - `OS10# show lldp neighbor`
10. After ensuring all the configurations are up and running, login to OMNI. From **OMNI Home** > SmartFabric instance > **Life Cycle Management** > **Replace Switch** to complete the switch replacement workflow.

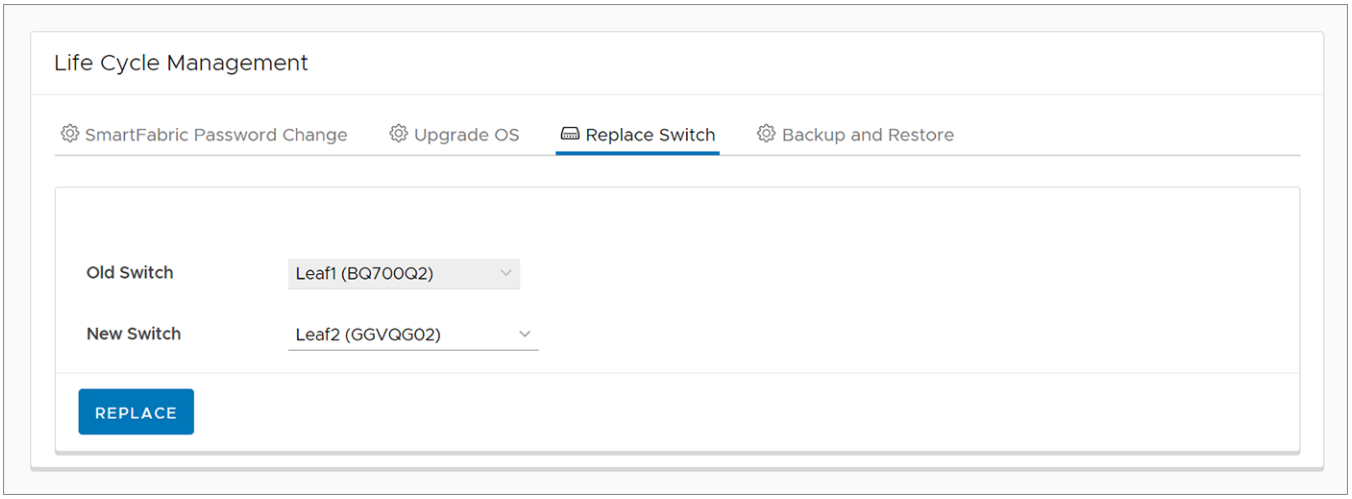

11. Select the switch that you want to replace from the list, select the new switch, and click **Replace**. The system displays switch replace success message.

Back up and restore the fabric configuration

You can save the current fabric configuration in a repository, and restore the data using a backup file when an error or failure occurs.

NOTE: This instruction is applicable for SmartFabric instance only and not supported on OME-M instance.

From OMNI, using the **Fabric backup and restore** feature, you can:

- Set a local or remote repository.
- Back up the configuration of a select fabric in the OMNI VM.
- Download the backup files to the local system.
- Delete the downloaded backup from the OMNI VM.

- Upload or import the fabric backup file from the local or remote repository to the OMNI VM.
- Restore the fabric from a backup file.

NOTE: The fabric backup and restore features are supported from the OS10.5.0.7 version. If the OS10 software version is less than 10.5.0.7, the system displays a message that backup is not supported for the software version and all the backup and restore functions are disabled.

Set Repository

To backup the configuration, set up a local repository on the OMNI VM or a remote repository to store the backup files. OMNI supports File Transfer Protocol (FTP) and Secure Copy protocol (SCP) to transfer the backup files to a remote repository.

NOTE: You can either set a local or a remote repository at a time. To change the backup repository, edit the repository setting accordingly.

Set a local repository

1. Select the SmartFabric Instance > **Life Cycle Management** > **Backup and Restore**.
2. From **Backup and Restore** tab, click **Set Repository**.

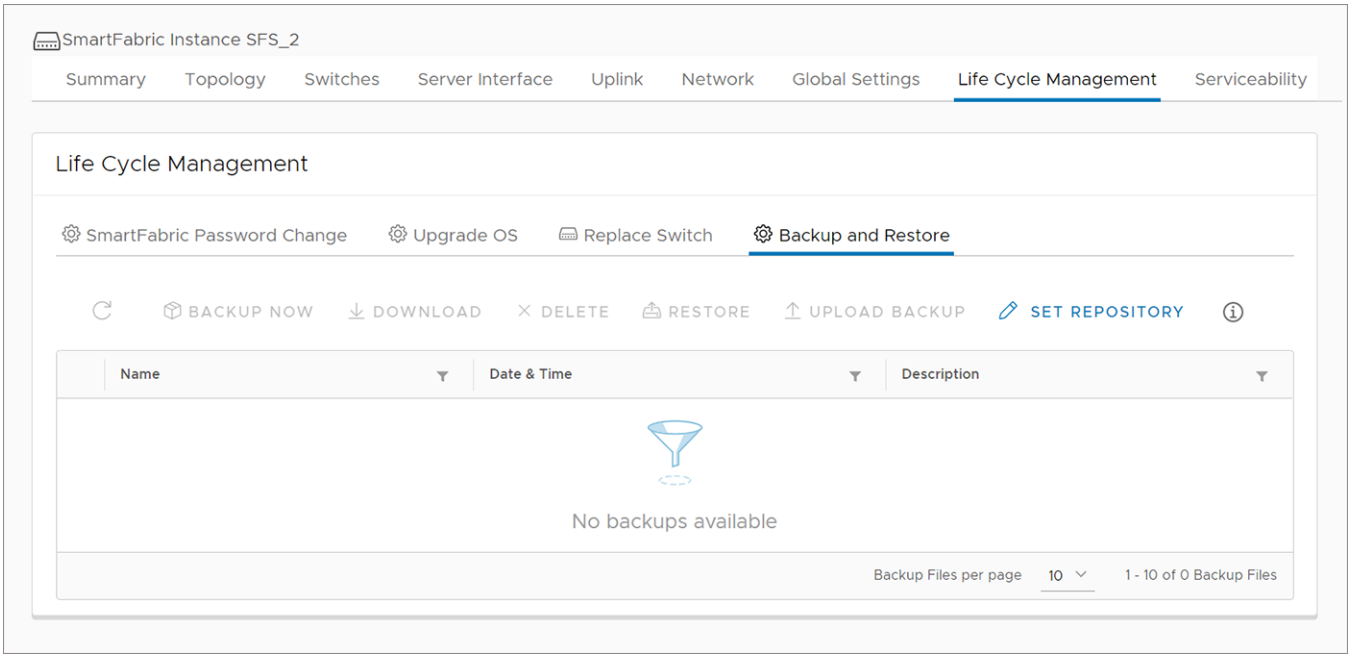

3. Select **Local**, and click **Submit**.

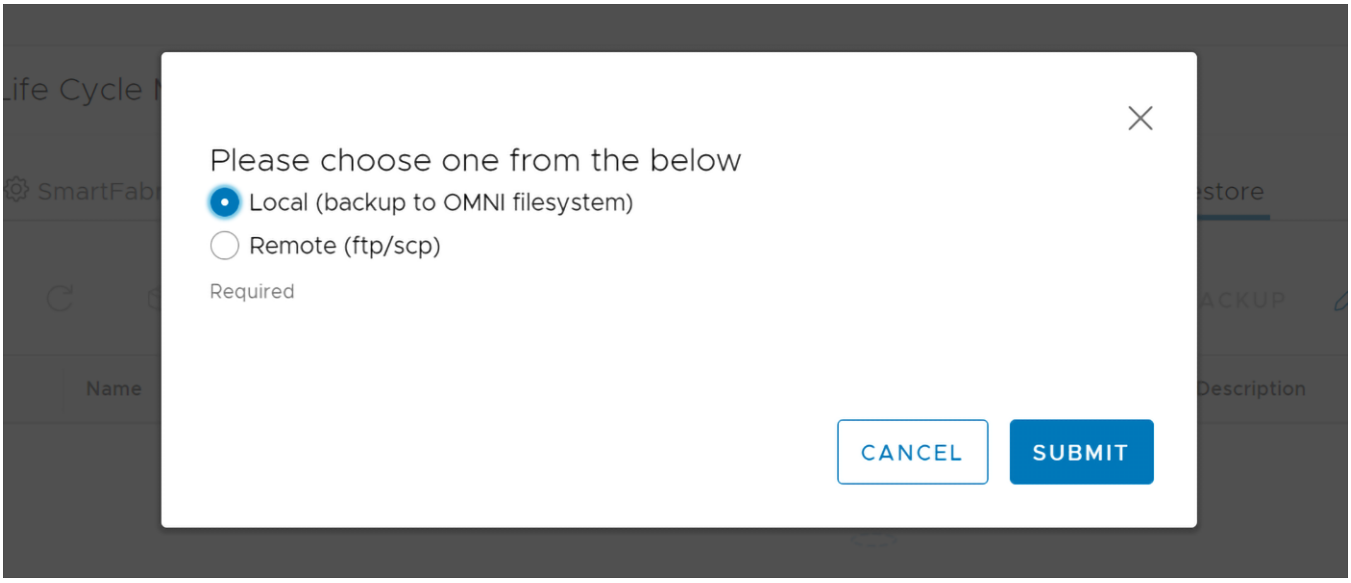

4. The system displays local repository configuration success message.

Set a remote repository

1. From **Backup and Restore** tab, click **Set Repository**.
2. Select **Remote**.
3. Select the protocol (SCP or FTP) from the list. Enter the **Hostname**, **Username**, and **Password** details.

   (Optional) Enter the **Repository Path** details, and click **Submit**.

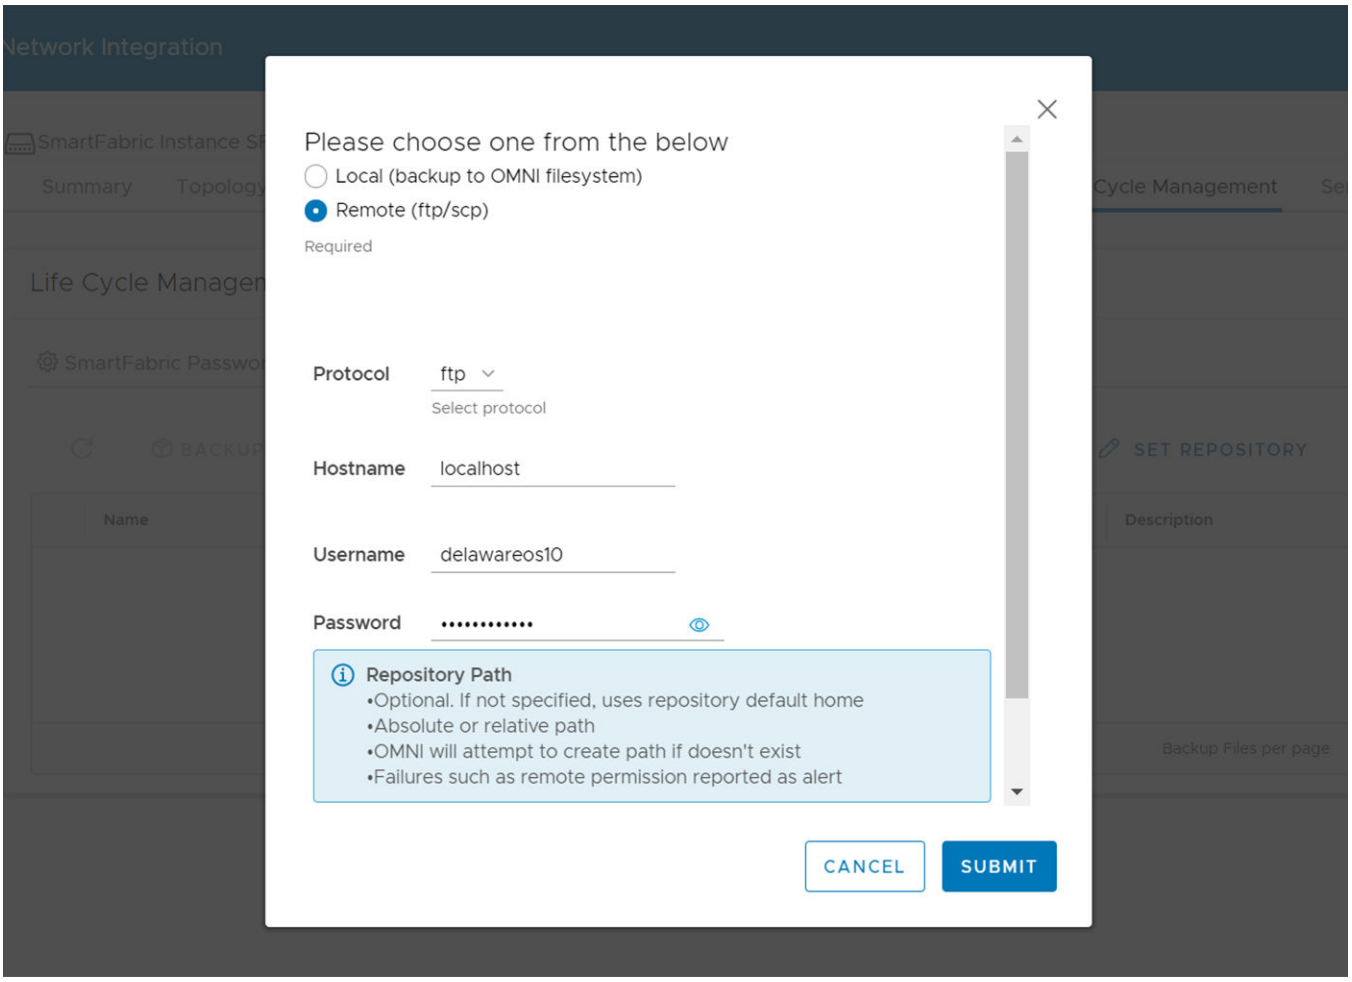

4. The system displays remote repository configuration success message.

View repository

View the repository details by clicking the information icon.

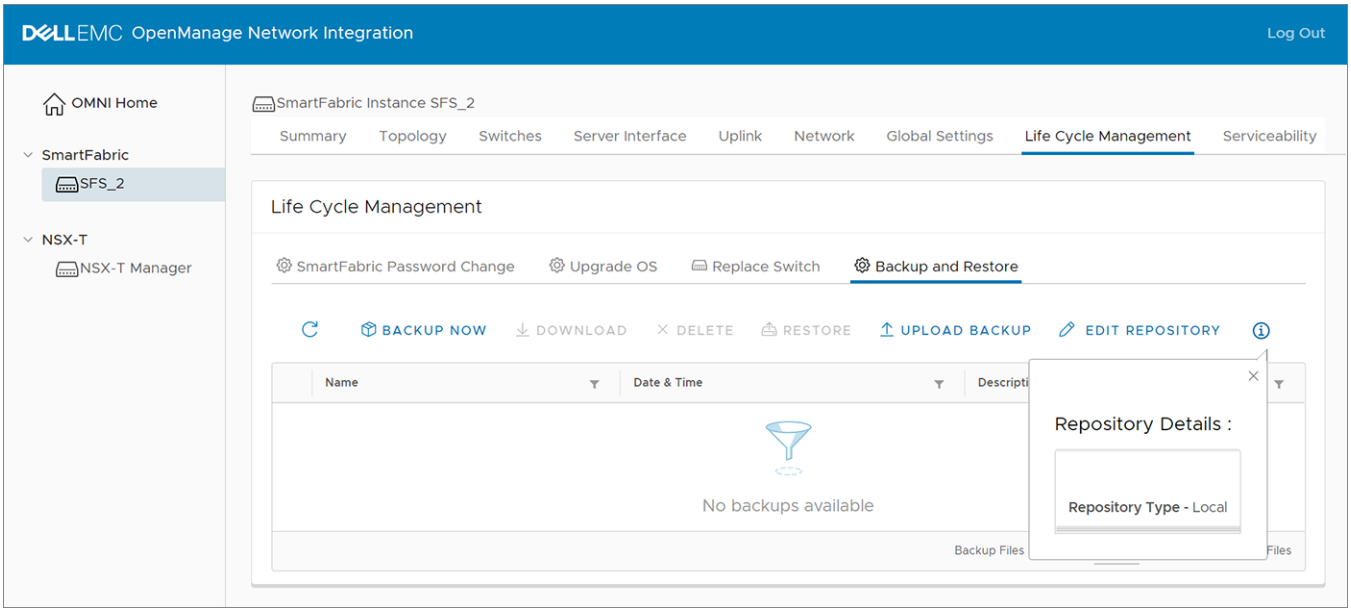

Edit repository

You can edit the repository type that is already set. To do so:

1. From **Backup and Restore** tab, click **Edit Repository**.
2. Edit the repository type, enter the required details if prompted, and click **Edit**.

   NOTE: When you edit the repository from local to remote, the backup files from the local OMNI VM are transferred to the remote repository. If you change the repository from remote to local, they backup files are not transferred to local OMNI VM.

Backup fabric configuration

To backup the fabric configuration:

1. Select **Life Cycle Management** > **Backup and Restore**, and click **Backup Now**.

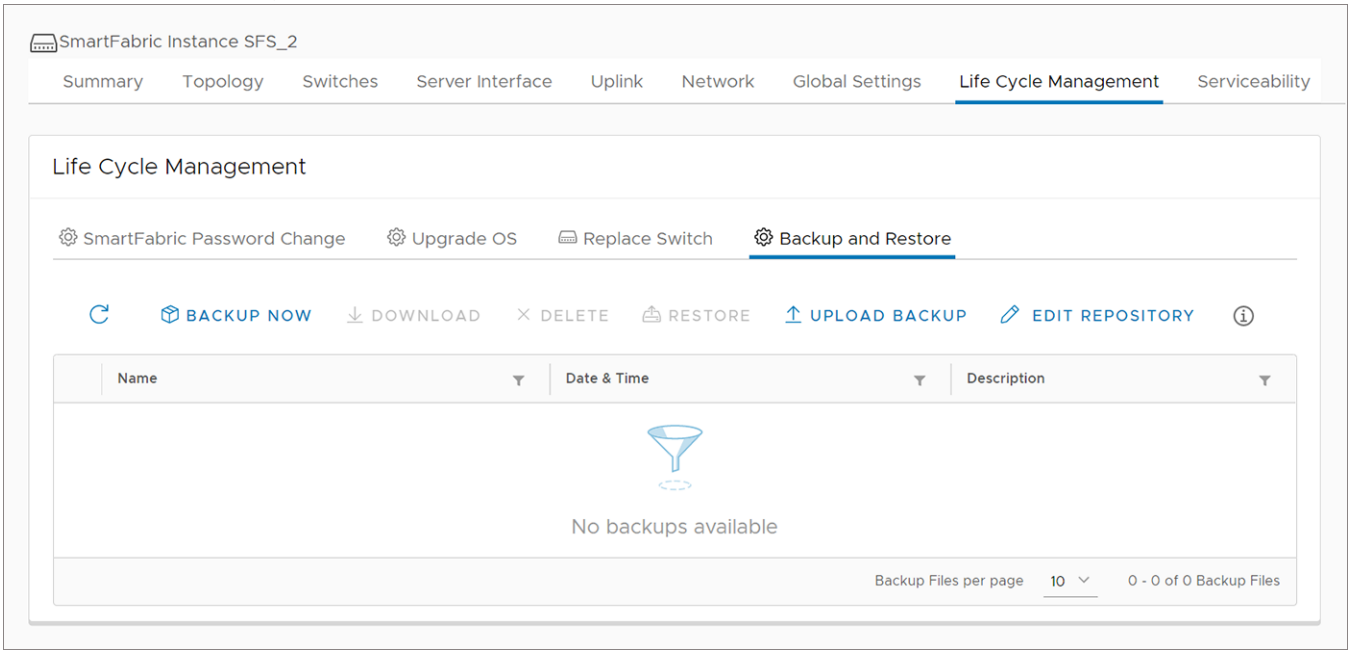

2. Enter the description for the backup file and click **Ok**.

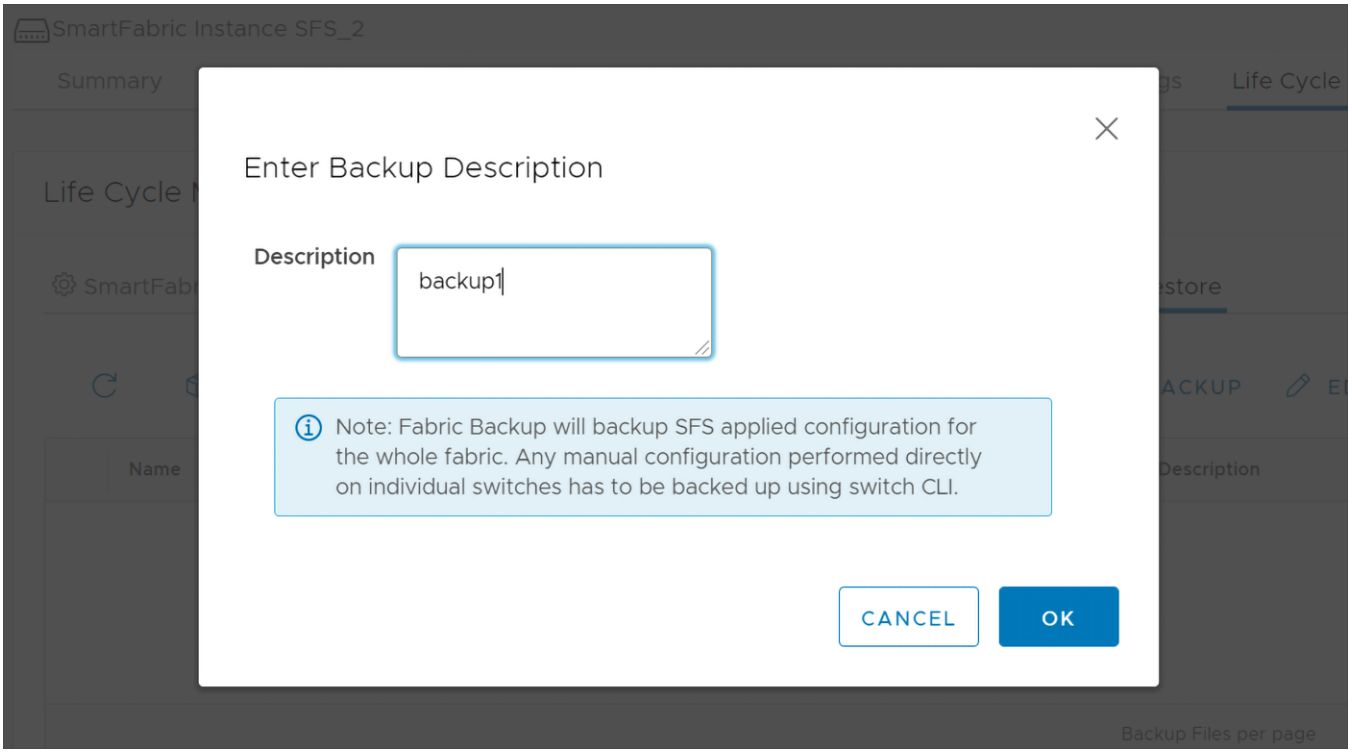

The backup file is stored as a JSON file.

NOTE: The backup action stores SFS-applied configuration for the whole fabric. Any OS10 system configuration that is done on the individual switches directly has to be backed up using the OS10 CLI. For more information about how to backup the configuration, see *Dell EMC SmartFabric OS10 User Guide*.

3. The system displays backup completed success message.

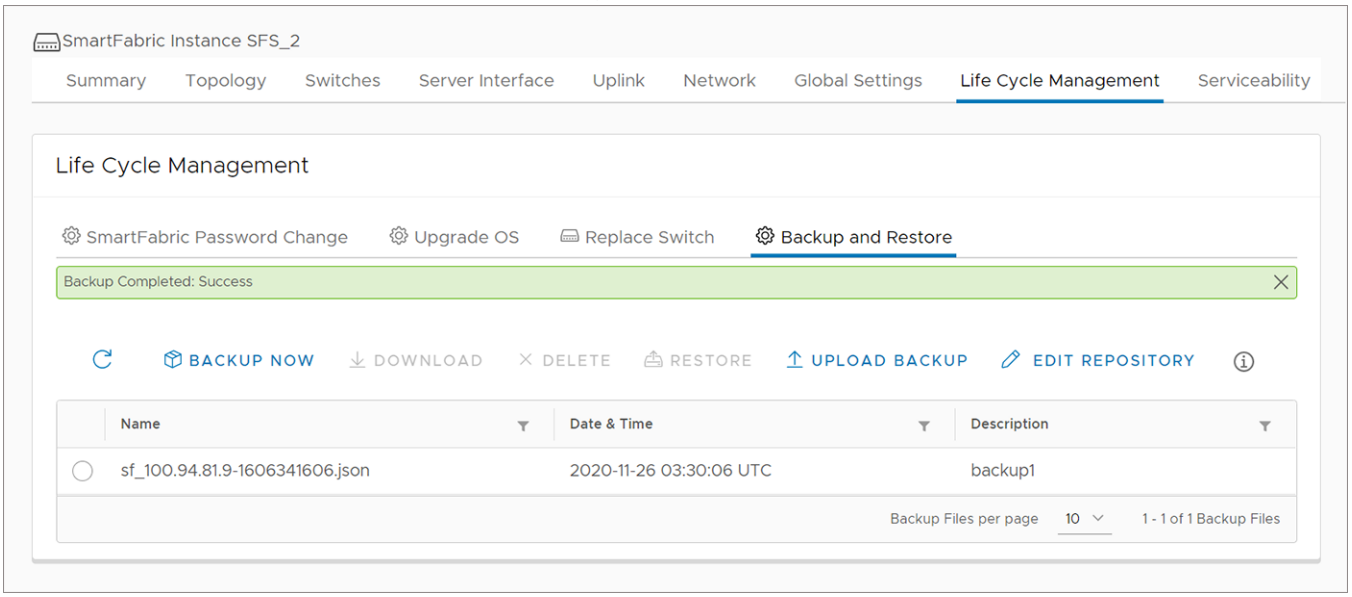

Download backup

You can download a backup file from the OMNI stand-alone VM to the local system. This option is available only when OMNI is accessed as a stand-alone application.

NOTE: **Download** option is not available when OMNI is launched as a plug-in from vCenter. Hence, you cannot download the backup JSON configuration files using OMNI plug-in.

1. Select **Backup and Restore** tab, and select the backup JSON file that you wanted to download from the list.
2. Click **Download**.

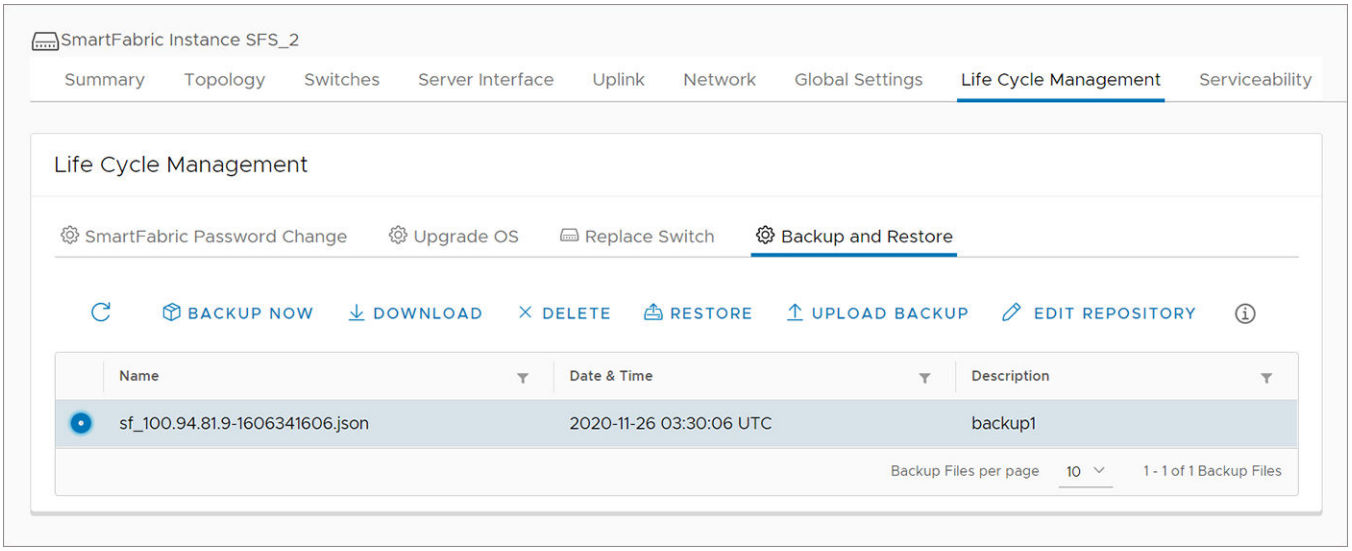

The file is downloaded locally with the backup download success message.

Delete backup

You can delete a backup file from the OMNI VM.

1. Select **Backup and Restore** tab.
2. Select the backup file that you want to delete from the displayed list, and click **Delete**.

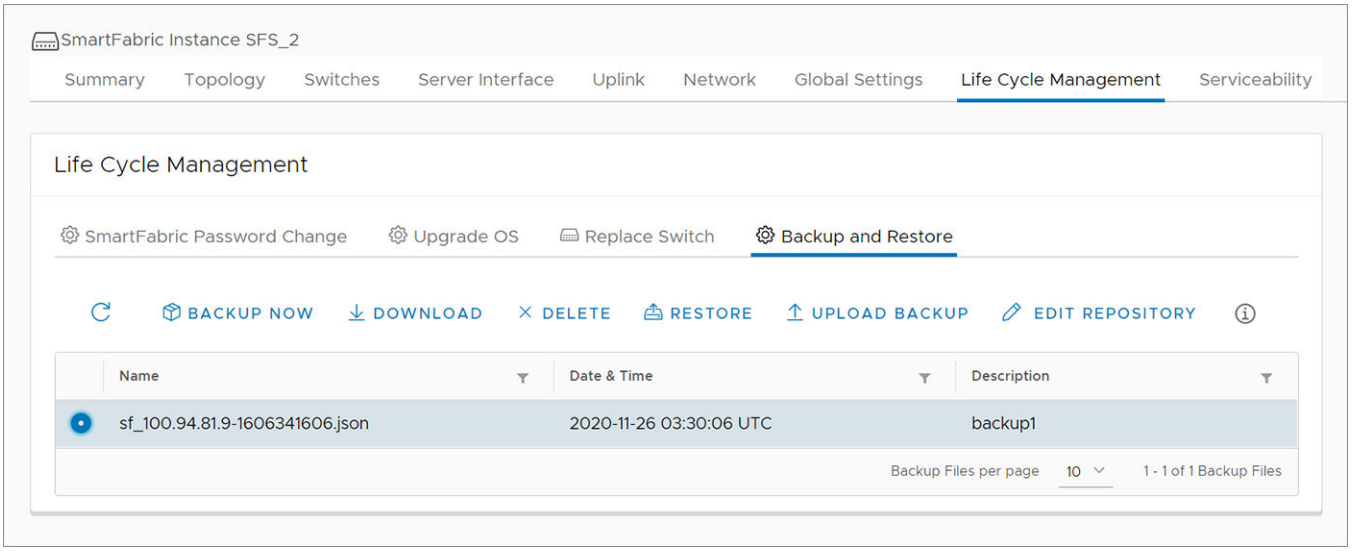

3. Click **Delete** to confirm.

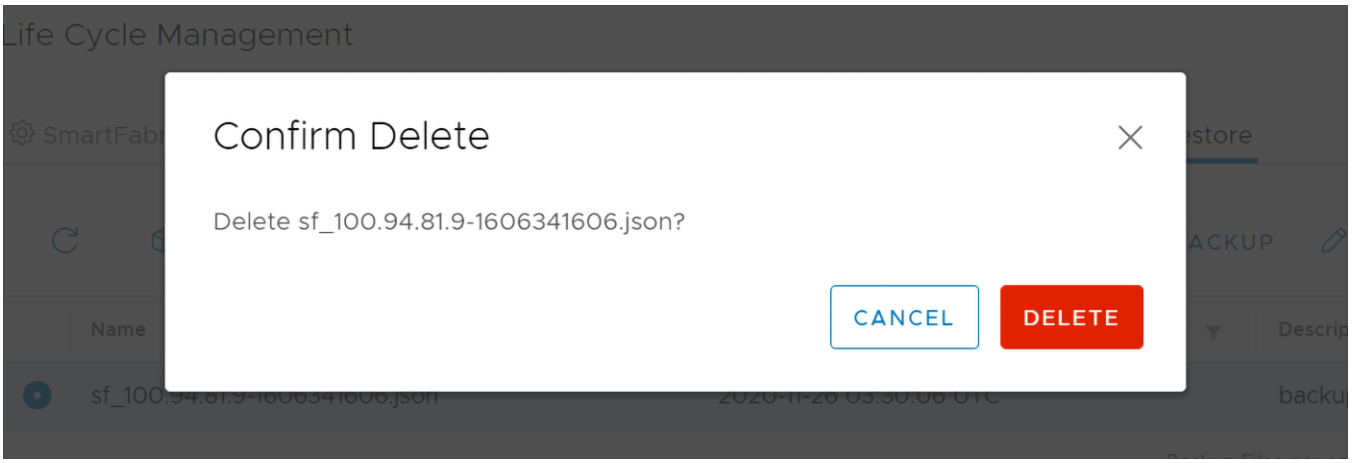

4. The system displays backup deleted success message.

Upload backup

You can upload a backup file from the local system to the OMNI VM.

1. From **Backup and Restore** tab, click **Upload Backup**.
2. Enter the description and choose the file that you want to upload, and click **Upload**.

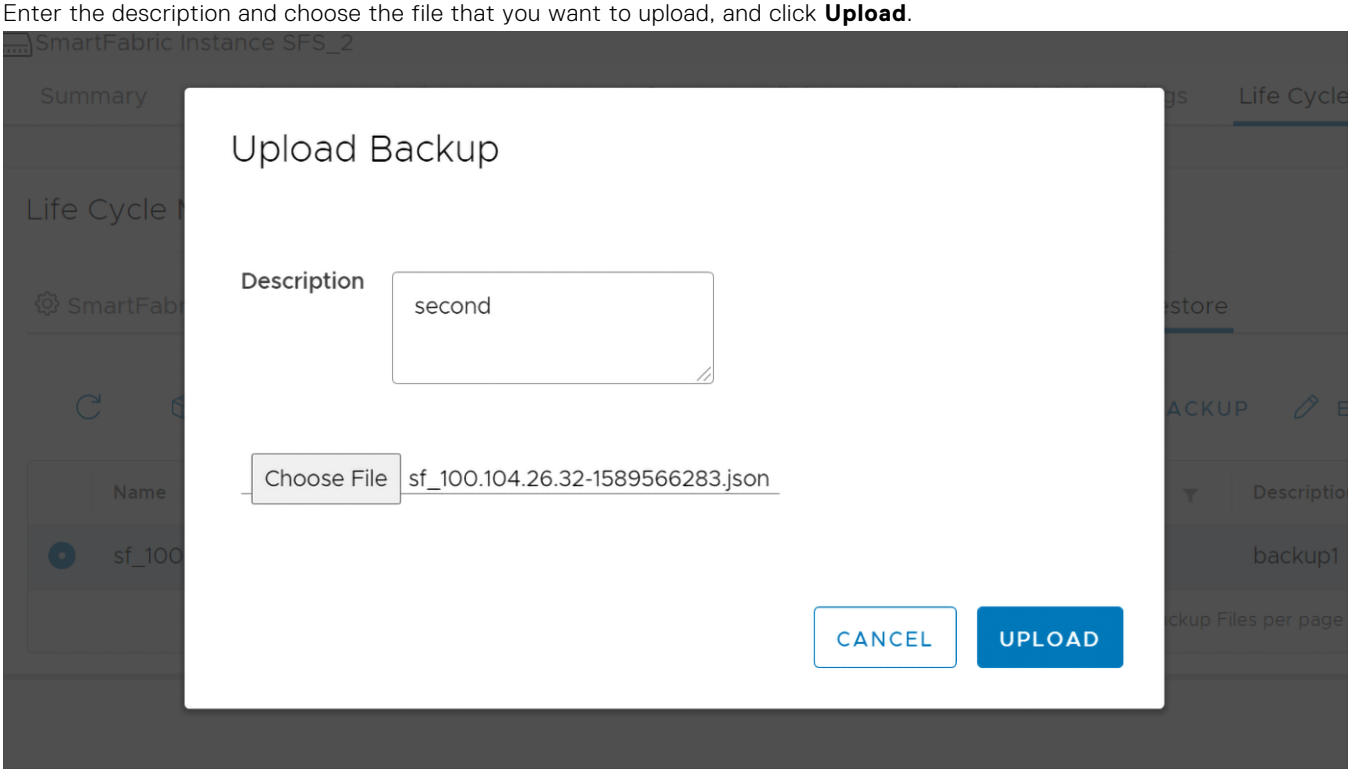

3. The system displays upload file success message.

NOTE: OMNI displays error if the uploaded file is not in the JSON format.

Restore from a backup file

You can restore the configuration running on the SmartFabric using a backup file during unexpected error situation or disaster.

CAUTION: Restore action is disruptive and cause connection downtime and traffic loss. The restore action erases all fabric configuration and restarts the entire fabric with the configuration in the backup file. It is highly recommended to use the restore action during a maintenance window.

1. Select **Life Cycle Management** > **Backup and Restore**.
2. Select the backup file from which you want to restore the configuration, and click **Restore**.

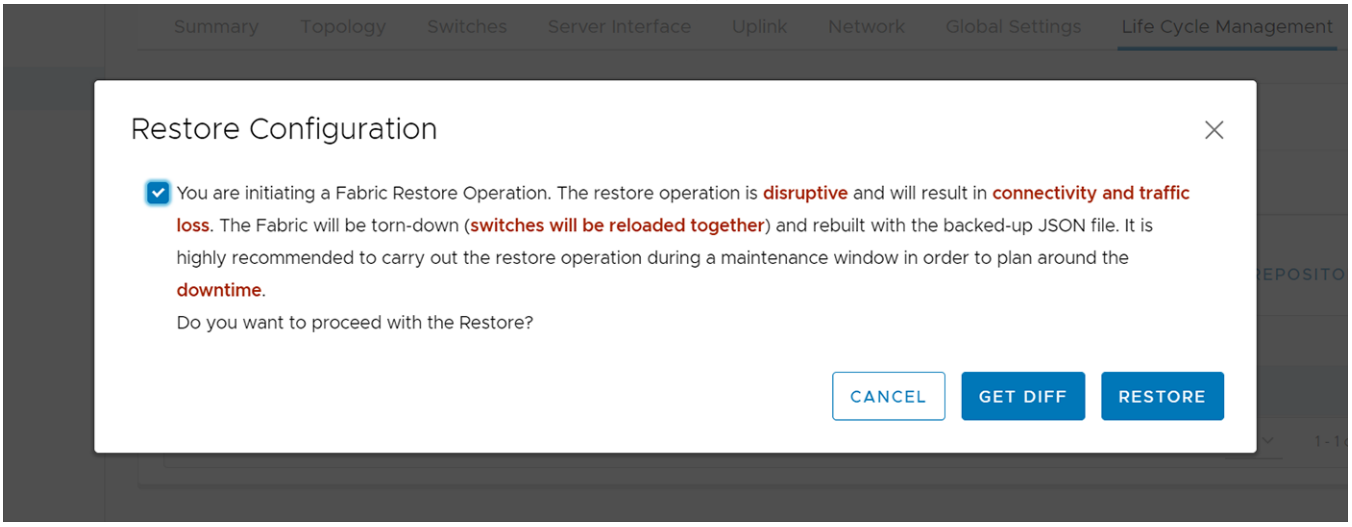

NOTE: The restore action reboots all the switches with the applied fabric settings. Any manual configuration that are performed directly on individual switches has to be restored manually using the OS10 CLI. For more information about how to restore the configuration, see *Dell EMC SmartFabric OS10 User Guide*.

3. (Optional) Click **Get Diff** to compare the current configuration with the configuration in the backup file.
Configuration Diff View displays the detailed comparison between the current and backup configuration.

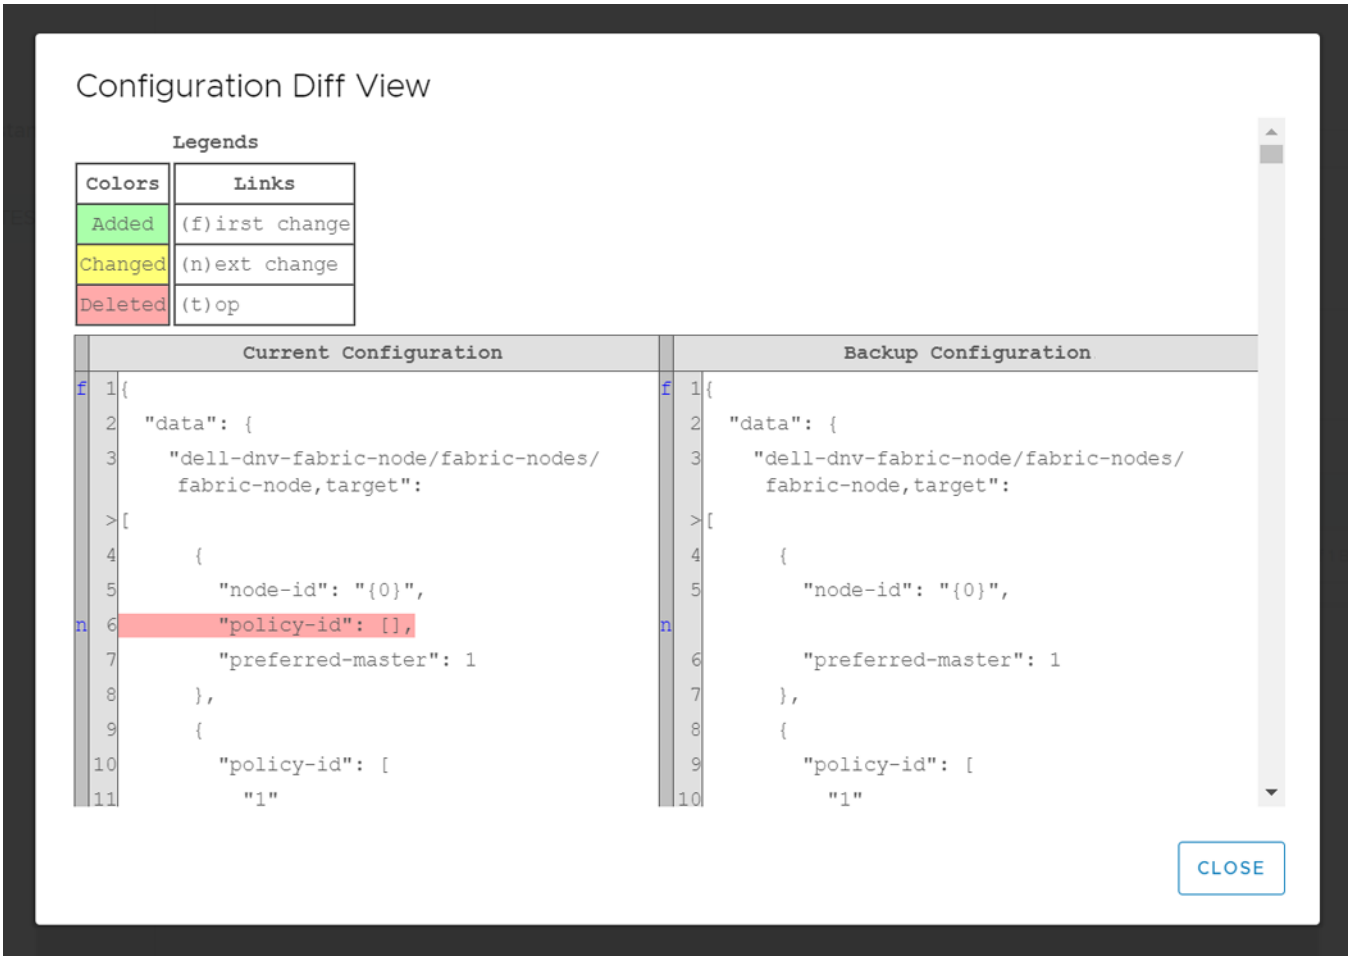

4. To proceed with the restore action, select the checkbox to confirm, and click **Restore**.

   Once you initiate the restore process, OMNI appliance changes the SmartFabric instance state to Maintenance mode automatically, which stops all the fabric automation services for that SmartFabric instance.

5. The system displays the restore success message.

   When the fabric restore is complete, change the Maintenance mode of the SmartFabric instance to **In Service**. For more information about Maintenance mode, see Maintenance mode.

6. For internal vCenter environment, restart the vCenter manually from the Platform Service Controller page. For more information about restarting the vCenter, see *VMware vSphere Documentation*.

11

Troubleshooting

Use the following information to troubleshoot some of the common problems that occur with the vCenter and OMNI appliance connectivity, OMNI UI launch, SmartFabric instance configurations, and OMNI automation.

Logs and support data for troubleshooting

You can generate support bundle with error and debug logs using OMNI. These logs can help to identify, diagnose, and debug problems.

Dell Technologies recommends downloading the support bundle from OMNI Appliance Management UI. By default, the log-level in OMNI appliance is set to ERROR. The appliance log can be swapped between ERROR to DEBUG. Change the log-level appropriately for each service and download the support bundle, see OMNI Appliance Management UI.

NOTE: Dell Technologies recommends setting the log level to DEBUG when you want to generate a support bundle during an issue.

If you cannot access the UI, use to OMNI console to download the support bundle. The support bundle is downloaded at `/tmp/support-bundle.tar.gz` on the OMNI VM. You can also change the log-level. When you change the log level from ERROR to DEBUG from OMNI VM console, the change applies to only the services `omni_api` and `omni_services`. For more information about OMNI management menu, see OMNI console menu.

Verify OMNI VM connectivity

After setting up OMNI, verify the IP address, DNS settings, and connection status from the OMNI VM console:

1. When OMNI is internal (deployed in one of the VxRail nodes in the cluster), ensure that you have configured IPv6 information for VxRail Mgmt network (ens192) and custom route as **fde1:53ba:e9a0:cccc::/64**. Disable IPv4 configuration for ens192 interface.
2. When OMNI is deployed on a ESXi server and registered with external vCenter, ensure that you have set IPv4 configuration with subnet mask and gateway information for vCenter server network (ens160). Set the IPv6 configuration for the interface to **Ignore**.
3. Check the interface connection status through OMNI console.

To view the interface configurations through OMNI console:

1. From the OMNI management menu, select **2. Interface Configuration Menu**.

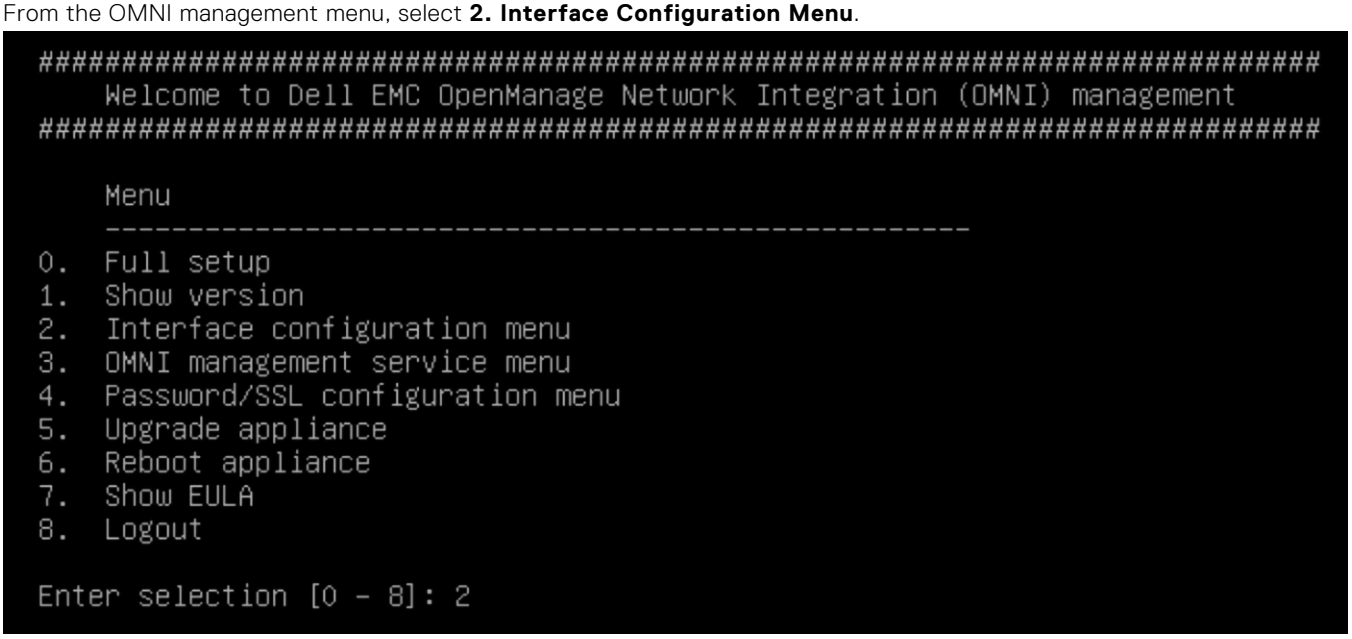

2. Enter the selection as **1. Show Interfaces** and press **Enter**.

```
                    ------------------------------------------------
                    OMNI interface configuration menu
                    ------------------------------------------------
                1.  Show interfaces
                2.  Show connection status
                3.  Configure interfaces
                4.  Show NTP status
                5.  Configure NTP server
                6.  Unconfigure NTP Server
                7.  Start NTP Server
                8.  Stop NTP Server
                9.  Exit

                Enter selection [1 - 9]: 1
```

```
Enter selection [1 - 9]: 1
sudo: unable to resolve host OMNI-1.3.14: Name or service not known
docker0: flags=4099<UP,BROADCAST,MULTICAST>  mtu 1500
        inet 172.17.0.1  netmask 255.255.0.0  broadcast 172.17.255.255
        ether 02:42:05:1a:45:da  txqueuelen 0  (Ethernet)
        RX packets 0  bytes 0 (0.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 0  bytes 0 (0.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

ens160: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 100.104.26.22  netmask 255.255.255.0  broadcast 100.104.26.255
        inet6 fe80::250:56ff:fe85:abb7  prefixlen 64  scopeid 0x20<link>
        ether 00:50:56:85:ab:b7  txqueuelen 1000  (Ethernet)
        RX packets 695002  bytes 159086623 (151.7 MiB)
        RX errors 0  dropped 54  overruns 0  frame 0
        TX packets 157180  bytes 144654105 (137.9 MiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

ens192: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        ether 00:50:56:85:93:cd  txqueuelen 1000  (Ethernet)
        RX packets 463229  bytes 46227842 (44.0 MiB)
        RX errors 0  dropped 52  overruns 0  frame 0
        TX packets 65686  bytes 11664357 (11.1 MiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
        inet 127.0.0.1  netmask 255.0.0.0
        inet6 ::1  prefixlen 128  scopeid 0x10<host>
        loop  txqueuelen 1000  (Local Loopback)
        RX packets 624296  bytes 90090468 (85.9 MiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 624296  bytes 90090468 (85.9 MiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

(END)
```

3. Select **2. Show Connection Status**. The status should be up and connected.

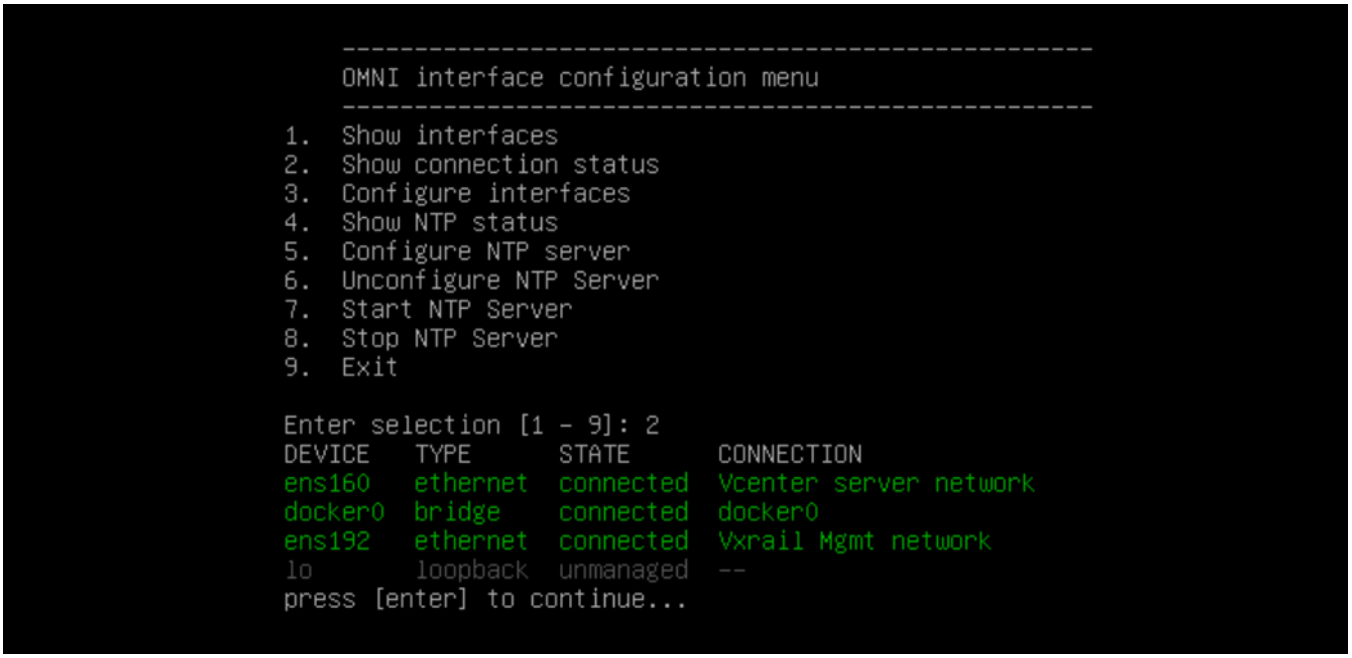

Unable to add SmartFabric instance in OMNI

Problem

Not able to add the SmartFabric instance in OMNI.

Causes

- SmartFabric instance is not reachable or down.
- IP address of the SmartFabric instance is not the master node IP address.

Resolution

- Ensure that the SmartFabric is reachable. To check the SmartFabric connectivity:
  1. Log in as **root** user through the OMNI VM console.
  2. Check the connectivity of the SmartFabric instance using the `ping` command. If OMNI is internal, use IPv6 address or hostname of the service instance. If OMNI is external, use IPv4 address or hostname of the service instance.

     ```
     :~$ ping <IPv4-address or hostname of the destination>
     :~$ ping6 <IPv6-address or hostname of the destination>
     ```

  3. If ping fails, check if the OMNI interfaces are configured properly, see Verify OMNI VM connectivity.
- Ensure that the IP address is the master node IP address. To check the IP address of master node:
  1. Identify the master node using the OS10 CLI command. For more information, see Add SmartFabric instance.
  2. Add the SmartFabric instance using the identified master IP address.

Missing networks on server interfaces

Problem

OMNI automation process fails to create and associate the appropriate network on a server interface during synchronization.

Causes

- If there is no relationship formed between the vCenter and service instances.
- Automation service is not running for that vCenter.

Resolution

- Check the relationship status between the vCenter and service instances. For more information, see Relationship information. If the relationship is not formed correctly, delete and reconfigure the SmartFabric instance and vCenter.
- If the southbound interface is a general ESXi server, create a server interface profile manually. If there is no server profile, no relationship is created. For more information about creating server interface profile, see server interface.
- Ensure that the service instance and vCenter are in **In Service** mode. The automation is not enabled if any of the relevant vCenters or SmartFabric instances is in **Maintenance** mode.

If the relationship status is correct and the automation is running, yet the issue persists, restart the automation service for the respective vCenter. Restart the automation services for the vCenter from the OMNI UI, see OMNI Appliance Management UI. After restart, OMNI synchronizes all the configurations again through automation.

Unable to launch OMNI UI

This information provides troubleshooting information when you are unable to launch OMNI plug-in from vCenter and as a stand-alone UI.

Unable to launch OMNI plug-in from vCenter

Problem

Unable to launch OMNI plug-in from vCenter.

vCenter does not show the OMNI plug-in option in the menu even after the vCenter is registered with OMNI through OMNI Fabric Management UI. vCenter also shows OMNI plug-in download errors, after the vCenter is registered with OMNI.

Causes

1. OMNI is not able to communicate with the vCenter due to SSL certificate errors.
2. vCenter could not resolve OMNI FQDN.

Resolution

1. Install a new SSL certificate, see Generate and Install SSL certificates.
   NOTE: After installing the certificate, unregister and re-register the vCenter instance again.
2. Ensure that the DNS is configured for the vCenter and is reachable. Also, DNS should have forward and reverse lookup configuration for OMNI FQDN or IP address.

If the problem still persists, try to unregister and register the OMNI appliance with vCenter again. For more information, see Register vCenter with OMNI.

Unable to launch stand-alone OMNI UI

Problem

Unable to launch the OMNI VM as a stand-alone application.

Cause

- vCenter server network connection (ens160) IPv6 configuration is not set to **Ignore**.
- OMNI essential services are not running.

Resolution

- Enter the IPv6 configuration for vCenter server network (ens160) to Ignore. For more information, see Setup OMNI.
- Check if the OMNI essential services are running using Appliance management UI. If OMNI UI is not accessible, check the OMNI management service status on the OMNI VM console. To check the services status:
  1. From the OMNI management menu, enter the selection as **3. OMNI management service menu**.
  2. Select **4. Restart OMNI management service** to restart all the database and web essential services.
     NOTE: To restart the automation services, go to OMNI Appliance Management UI and restart the services.

3. Select **2. View OMNI management service status** to view the list of registered vCenter managed by the OMNI VM. Confirm that all services are active.

```
Enter selection [1 - 7]: 2
           Name                          Command             State              Ports
-------------------------------------------------------------------------------------------------
omni_api                        /bin/bash -c python -c "fr      Up
                                ...
omni_api_celery_worker          celery worker --app=vcente      Up
                                ...
omni_automation_app_celery_be   celery beat --app=vcentera      Up
at                              ...
omni_automation_app_celery_wo   celery worker --app=vcente      Up
rker                            ...
omni_db                         docker-entrypoint.sh postg      Up      127.0.0.1:5432->5432/tcp
                                ...
omni_events_celery_beat         celery beat --app=vcentera      Up
                                ...
omni_events_celery_worker       celery worker --app=vcente      Up
                                ...
omni_events_receiver            /usr/local/bin/gunicorn -w      Up
                                ...
omni_nginx                      nginx -g daemon off;            Up
omni_queue                      docker-entrypoint.sh rabbi      Up      15671/tcp,
                                ...                                     127.0.0.1:15672->15672/tcp,
                                                                        15691/tcp, 15692/tcp,
                                                                        25672/tcp, 4369/tcp,
                                                                        5671/tcp,
                                                                        127.0.0.1:5672->5672/tcp
omni_services                   /bin/bash -c python -c "fr      Up
                                ...
omni_services_celery_worker     celery worker --app=vcente      Up
                                ...
/usr/local/lib/python3.5/site-packages/OpenSSL/crypto.py:12: CryptographyDeprecationWarning: Python
3.5 support will be dropped in the next release of cryptography. Please upgrade your Python.
  from cryptography import x509
2020-12-04 08:36:43,387 OMNI is registered with 100.104.26.63 vCenter host
press [enter] to continue...
```

NOTE: View OMNI management service status is recommended for status validation and debugging purpose. Hence, the output does not show the port numbers.

OMNI plug-in does not show service instance

Problem

OMNI plug-in does not show any service instance even though the service instance is added to OMNI.

OMNI plug-in does not show the instance when the vCenter is launched using the IP address but the vCenter is registered with FQDN in OMNI.

Cause

This problem can happen when the DNS is either not reachable or not configured with required settings.

Resolution

Ensure that the DNS is reachable, and configured with forward and reverse lookup for vCenter IP address or FQDN.

Unable to register the vCenter in OMNI

Problem

Unable to register the vCenter in OMNI.

Causes

- vCenter server network (ens160) is not assigned with correct port-group during deployment.
- IP addresses assigned to interfaces of OMNI are on the docker private network (172.16.0.0/16).

Resolution

- Ensure that ens160 is connected to the vCenter server network properly during OMNI deployment. For more information, see Setup OMNI.
- Change the docker private network configuration, see Configure docker private network.

OMNI is unable to communicate with other devices

Problem

OMNI appliance is unable to communicate with any external devices.

Causes

When there is a conflict between the default OMNI docker private network and any other network to which OMNI is connected, OMNI cannot communicate with the devices in that network.

The conflicts occur when:

- The ens160 and ens192 interfaces have IP addresses assigned from the docker private network (172.16.0.0/16).
- Any external entity such as vCenter, SFS instance, OME-Modular, NSX-T, NTP server, DNS server, and so on, which has IP address that is assigned from the docker private network.
- OMNI is connected to a larger network in which one or more subnetworks IP range overlaps with the docker private network.

Resolution

Change the docker private network configuration, see Configure docker private network.

NOTE: OMNI appliance reboots after the docker private network IP address is changed.

Timestamp not synchronized in OMNI

Problem

Logs and events timestamp details are not synchronized with the current data center.

Cause

OMNI does not have the proper NTP server configuration.

Resolution

Check the NTP server configuration in the OMNI appliance. Apply the correct configuration, if required. To check and change the NTP server setting:

1. From the OMNI management menu, select **2. Interface Configuration Menu**.

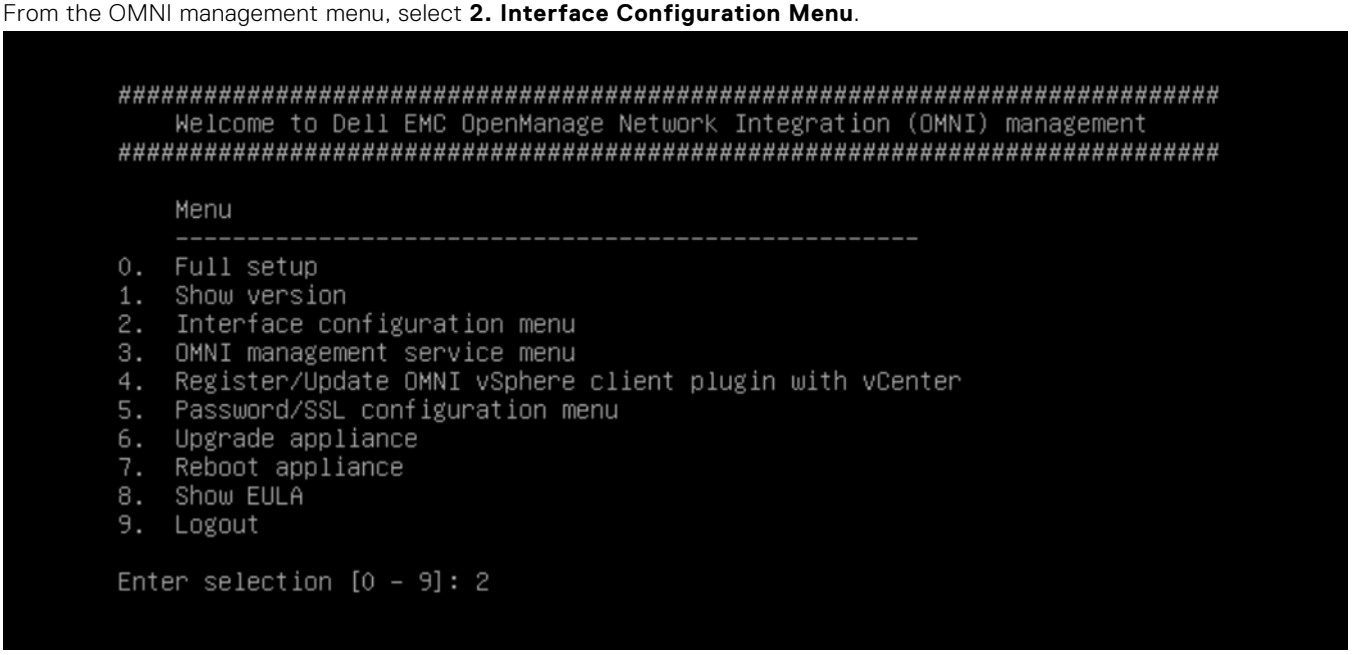

2. Select **4. Show NTP Status**.

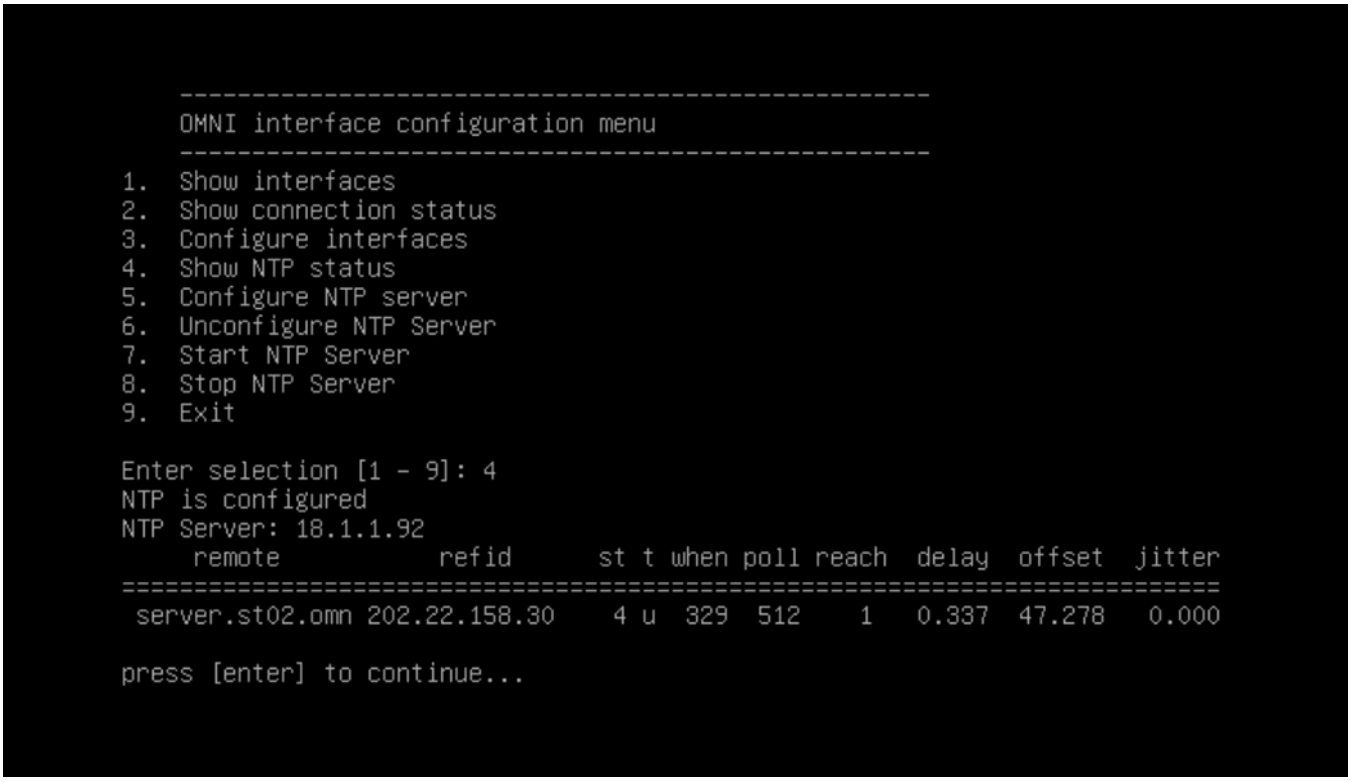

3. If the NTP server is not configured to the correct data center, select **5. Configure NTP Server**, and enter the valid NTP server IP address or hostname.

Reset OMNI VM password

To change the log-level from OMNI console:

1. Reboot the VM from vCenter, then select **Advanced Options for Debian GNU/Linux**.

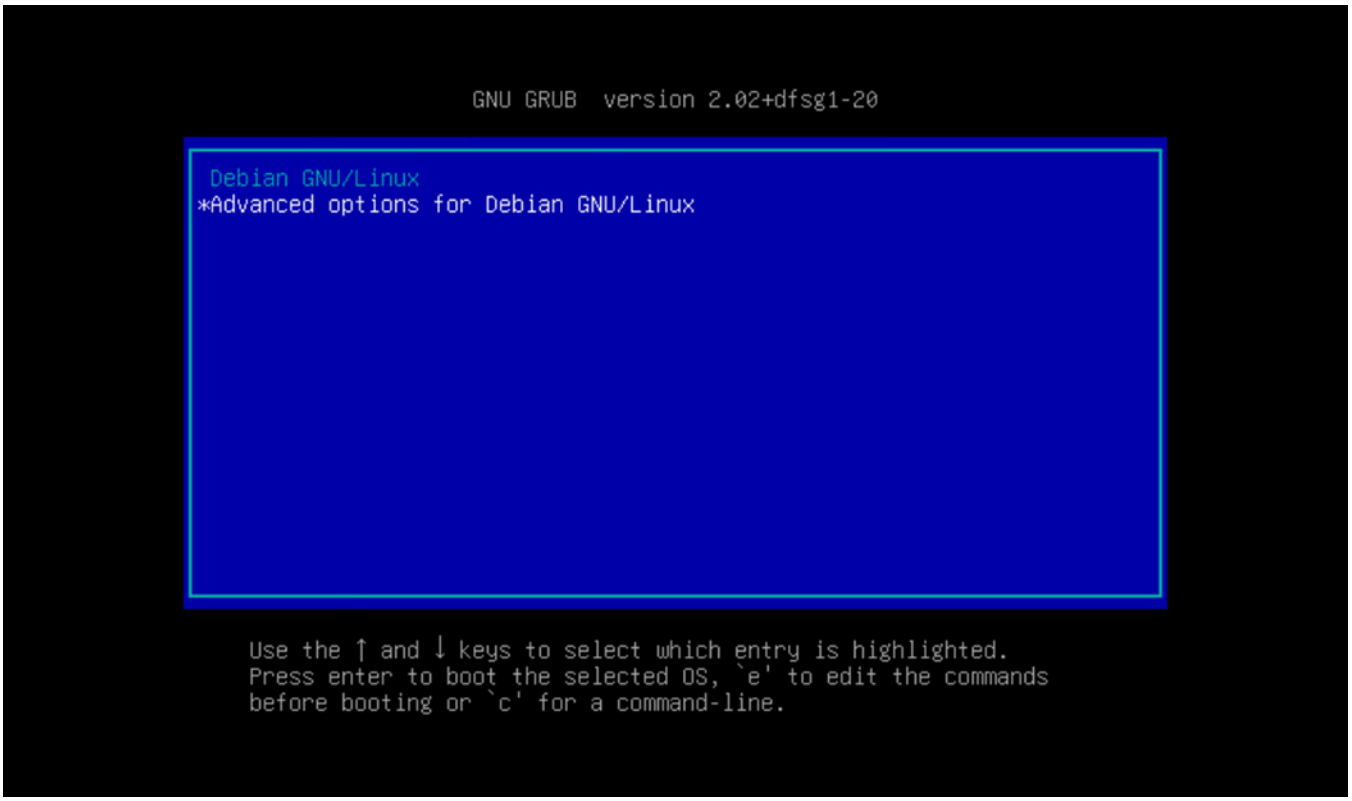

2. Use the arrow keys to go to the line starting with `linux` and ending with `ro quiet`.

3. Append `init=bin/bash` after `ro quiet`.

4. Press **Ctrl-X** to boot into the shell with root access.

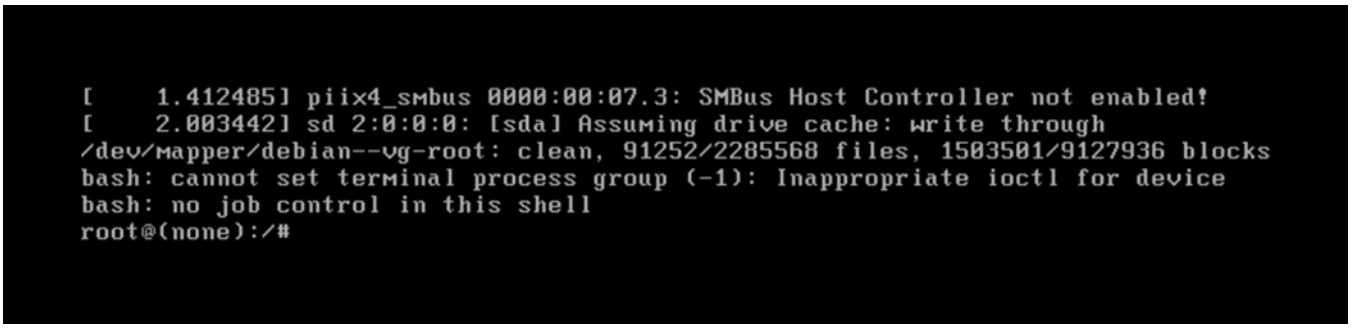

5. Remount the directory.

```
# mount / -rw -o remount
```

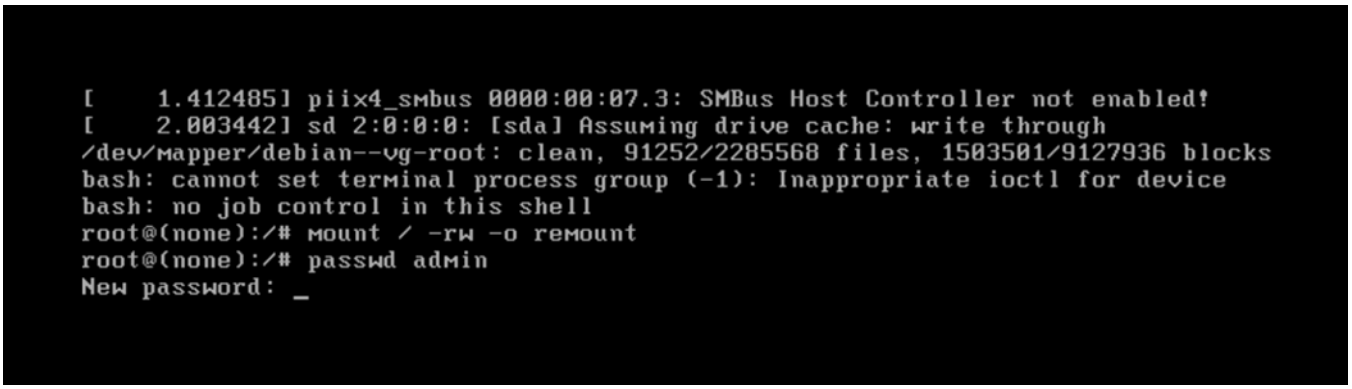

6. Change the password for admin using `passwd admin`. Enter the new password and confirm the password.

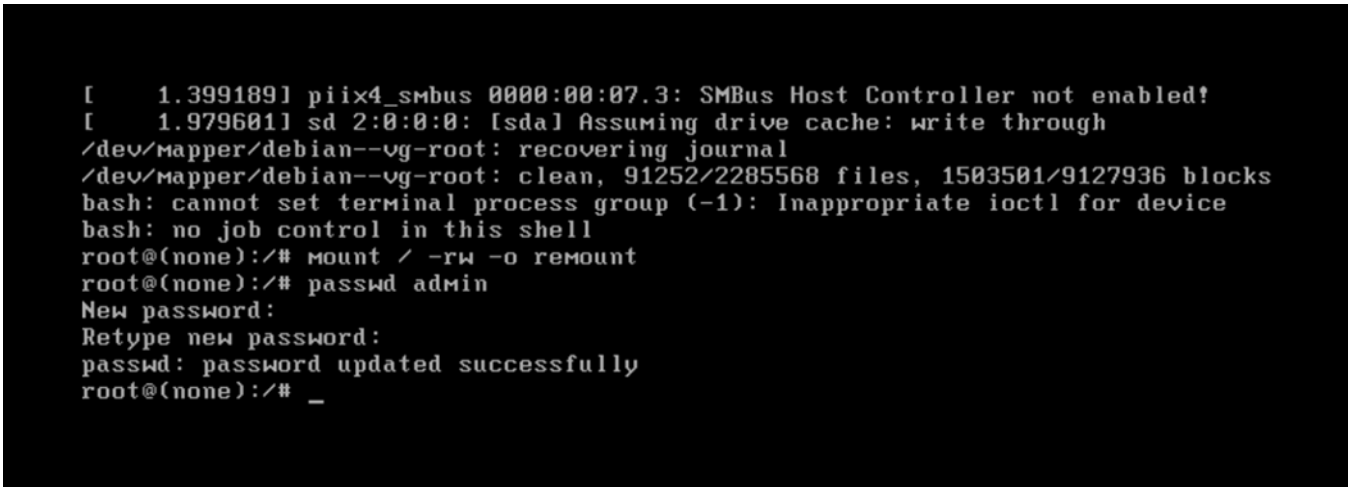

7. Reset the VM from vCenter and log in through the new password for the OMNI VM.