SmartFabric Services for OpenManage Network Integration User Guide

Release 1.2

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.

2020 - 2

Rev. A00

Contents

1

SmartFabric vCenter

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) are an automation framework that is built into the Dell EMC SmartFabric OS10 network operating system, integrating 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 customers existing data center networks.

OpenManage Network Integration

Dell EMC OpenManage Network Integration (OMNI) is a management application that is designed to complement SFS, providing a web-based GUI 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
- Deployed as an external plug-in for VMware vCenter. OMNI when deployed as a plug-in for VMware vCenter enables:
  - Enables zero-touch automation of physical underlay network fabric running SFS corresponding to changes in the virtual network layer
  - Extends vCenter Host Network Inventory data to include physical switch connectivity details for easy monitoring and troubleshooting
  - Enables single pane of management for one of more SmartFabric instances through the OMNI portal pages that are embedded within vCenter

VxRail SFS integration solution

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 to automate 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/vCenter
- Better infrastructure visibility
  - Tight integration between VxRail appliance and Dell EMC ON-Series PowerSwitches
  - Fabric connectivity extended to PowerSwitches required for VxRail solutions only
- Improved SLA

    - Fully validated software stack recommendation
    - 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
    - **Leaf/ToR switches**: 10 GbE — S4112F-ON, S4112T-ON, S4128F-ON, S4128T-ON, S4148F-ON, S4148T-ON; 25 GbE — S5212F-ON, S5224F-ON, S5248F-ON, and S5296F-ON
    - **Spine switches**: S5232F-ON and Z9264F-ON
- 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 Dell EMC VxRail Support Matrix for the latest software releases that support the VxRail SmartFabric Service-integrated solution. For complete information about deploying a VxRail SmartFabric solution, see Dell EMC VxRail Networking Solutions.

More resources

- Dell EMC SmartFabric OS10 User Guide, Release 10.5.0
- Dell EMC VxRail Appliance Administration Guide, Release 4.7.x
- Dell EMC VxRail Appliance Software 4.7.x Release Notes
- Dell EMC VxRail Multirack Deployment Guide
- Dell EMC VxRail QuickStart Guide
- Dell EMC VxRail Network Planning Guide

2

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 network operating system 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 PowerEdge modular system switches (PowerEdge MX), and PowerSwitch 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.

For complete information about SFS for PowerEdge MX fabric, see Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide.

Topics:

SFS for data center 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/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

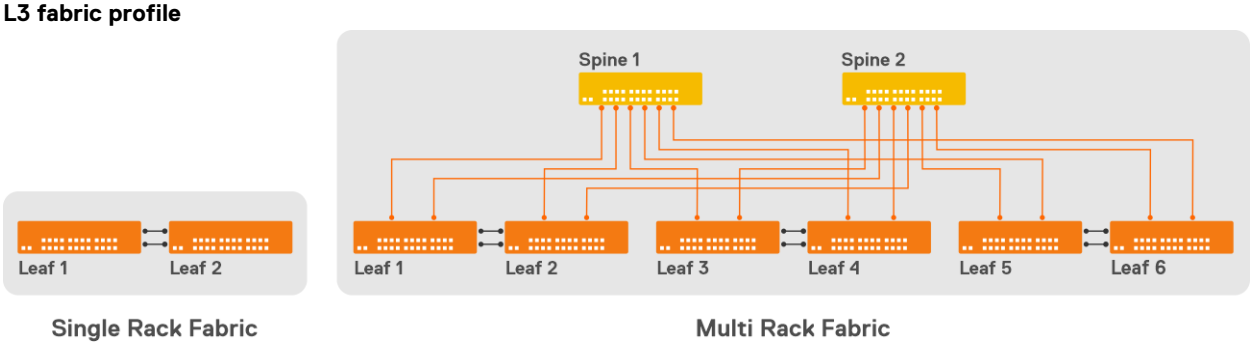

Single Rack Fabric

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 Dell EMC VxRail 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, Release 10.5.0.

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 specified before enabling SFS.

NOTE: The VLTi ports (ICL ports) cannot be modified once SFS in enabled. It is recommended to select the required number of ports upfront. SFS must be disabled and reenabled again to change the VLTi ports which can result in service interruption.

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 1/1/1-1/1/5
  ```

- 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, Release 10.5.0.

SFS Graphical User Interface

You can also enable SFS using the SFS Graphical User Interface (GUI). OS10 switches support SFS GUI to set up initial SFS configuration in SFS leaf and spine deployment. The SFS GUI 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 GUI, see SmartFabric Services in the *Dell EMC SmartFabric OS10 User Guide, Release 10.5.0*.

Fabric creation

This information describes switch discovery, SFS Master, Master advertisement, SFS REST services, Master high availability, preferred Master, SFS domain or network fabric, rack/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 Fabric 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 (ICL), and the port where another spine switch is discovered is configured as an interswitch link (ISL). A switch operating as a spine will 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 Keepalived 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 day 2 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 it is recommended to change this 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 are always had 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

You can select a designated set of leaf switches as the 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 other leaf switch can become the Master.

When fabric is expanded, newly added switches may come up and from a fabric among themselves, and elect a Master before they are connected to the existing fabric. When the new fabric merges with the running fabric, Keepalived protocol elects a Master switch from the new leaf switches joining the fabric which overwrites the configuration in the existing fabric. It is critical to ensure that the leaf nodes in the existing fabric are set up to be Preferred Master before expanding the fabric to prevent this.

NOTE: When you create an uplink using the SFS GUI or OMNI detects at least one uplink, the Preferred Master is automatically set on all leaf switches in the fabric, then.

SFS domain or network fabric

SFS domain or network fabric is interchangeable terminology, and the fabric consists of all switches directly connected to form a single logical network fabric. 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/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, the 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 from 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 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.

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* specifically 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. Only RPVST+ mode is supported on switches in SFS mode to have nonblocking leaf and spine connections. To prevent the loops in SFS, control VLAN 4000 RPVST is enabled with root bridge that is forced to be on one spine switch.

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.

Spanning-tree protocol is disabled for VxLAN networks. SFS automatically creates user networks as VxLAN networks inside the fabric and these networks. For 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.

For Layer 3 uplinks using routed Interfaces, spanning-tree is disabled on the uplink ports automatically. For Layer 3 uplinks using VLAN IP interfaces, RPVST+ is enabled on the VLAN and cannot be disabled or changed.

SFS VxRail integrated solutions

This information describes two SFS VxRail integrated solutions.

- **SFS for VxRail L2 single rack** — 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 L3 leaf and spine fabric personality for future expansion.
- **SFS leaf and spine fabric** — 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 leaf and spine fabric
- SFS VxRail L2 single rack fabric

Table 1. VxRail SFS personality comparison

VxRail L2 single rack fabric	Multi rack leaf and spine fabric
Network fabric with two ToR switches in a single rack, and 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 VxRail 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 prior to SmartFabric OS10 10.5.0.5.	All new VxRail 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 just role and VLTi ports for leaf as fixed parameters. Enable SFS using SmartFabric GUI also. For more information about SFS GUI, 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. These are fully customizable as part of Day 2 operations.

VxRail L2 single rack fabric	Multi rack leaf and spine fabric
All networks created during initialization, VxRail deployment and Day 2 operations are VLAN backed network with customer router acting as the gateway.	Networks that are created during initialization and the ones created as part of VxRail deployment and vCenter integration are VxLAN 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, and will be available in a future release.

NOTE: We recommend that all new deployments for the VxRail integration solution be enabled with leaf and spine fabric for single rack or multi rack deployments. VxRail SmartFabric deployments using older VxRail L2 single rack fabric cannot be upgraded to the new leaf and spine fabric automatically. A migration workflow will be available in a future release to allow existing deployments to expand to a multi rack solution.

3

OpenManage Network Integration

OpenManage Network Integration (OMNI) is a component 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 SFS 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 2. 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 2 virtual CPUs; 2 GB memory; 40 GB hard disk	Out-of-band (OOB) management network Provides reachability to DNS, default gateway, and where OMNI obtains the IP/hostname Provides reachability to Management network (vCenter IP/hostname, SmartFabric Management IP/hostname) VxRail default: vCenter Server network	In-band link-local network Provides reachability to SmartFabric link-local network for IPv6 VIP reachability VxRail default: VxRail Management network	vCenter HTML5 (/ui) plug-in; click OpenManage Network Integration link OMNI stand-alone user interface: https://*OMNI_IP_or_hostname*/delawareos10/ using `admin` user SSH to OMNI VM IP/hostname as `admin` user OMNI VM console using vCenter/ESXi `admin` or `root` user

NOTE: Even when OMNI is deployed in-band, it is recommended to set 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.

Topics:

- OMNI virtual appliance creation
- OMNI setup
- OMNI vCenter client plug-in registration

OMNI virtual appliance creation

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

NOTE: The OMNI portal or SmartFabric Services user interface does not provide localization.

Download and install OVA

1. Download the OVA from OpenManage Network Integration support, then store the OVA image locally.
2. In the vSphere Client, select **Hosts and Clusters**, right-click the cluster that the plug-in must manage, then select **Deploy OVF Template**.

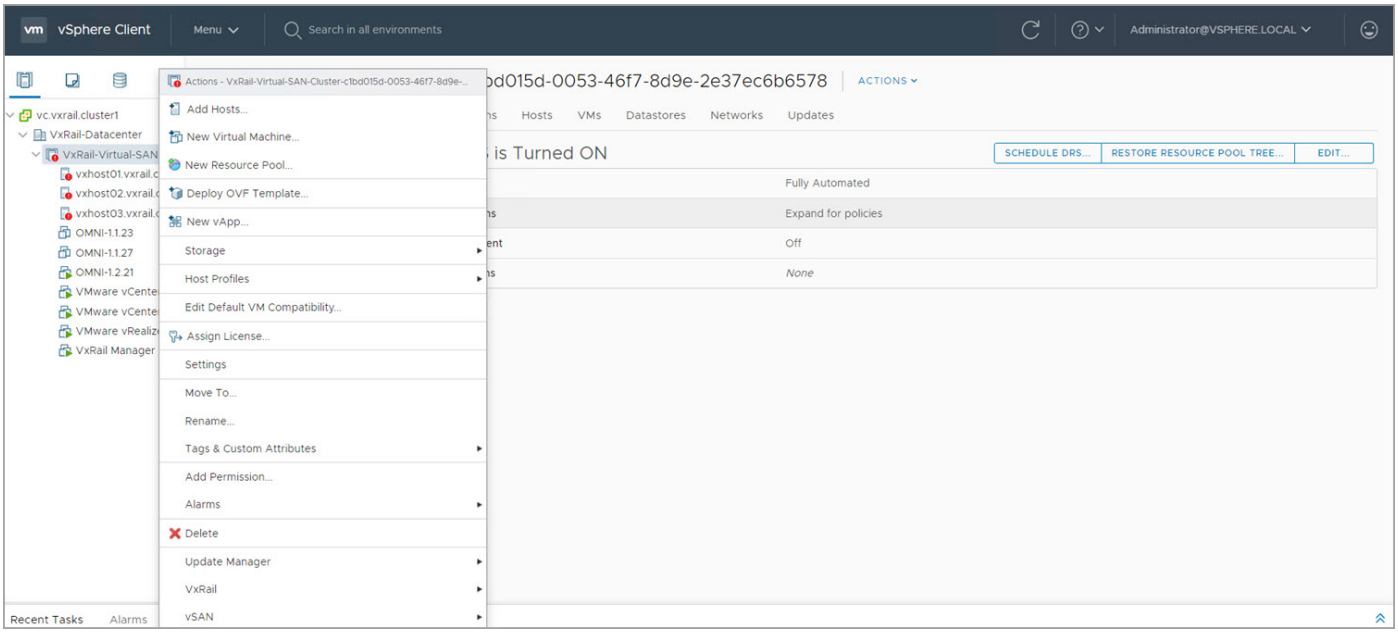

3. Select **Local file**, click **Choose Files** and select the OMNI ova file from a local source, then click **Next**.

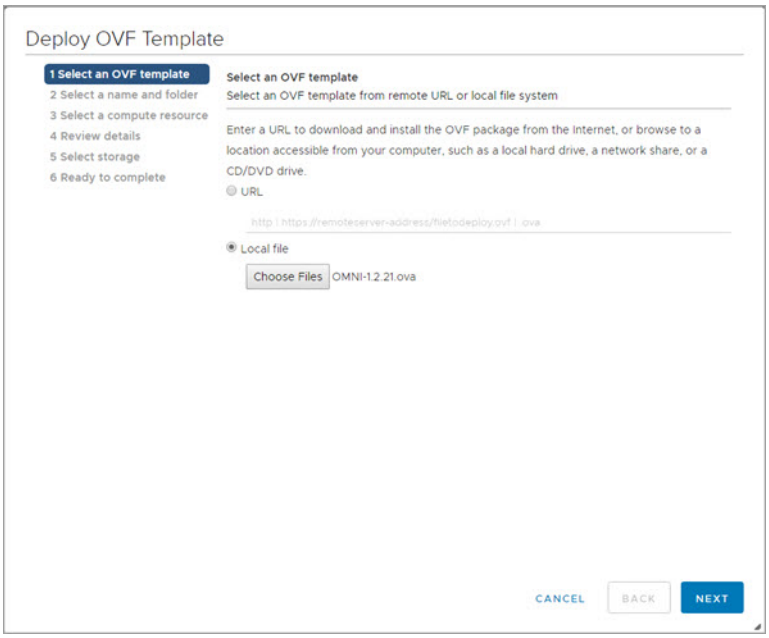

4. Select a name and folder for the VM, then click **Next**.

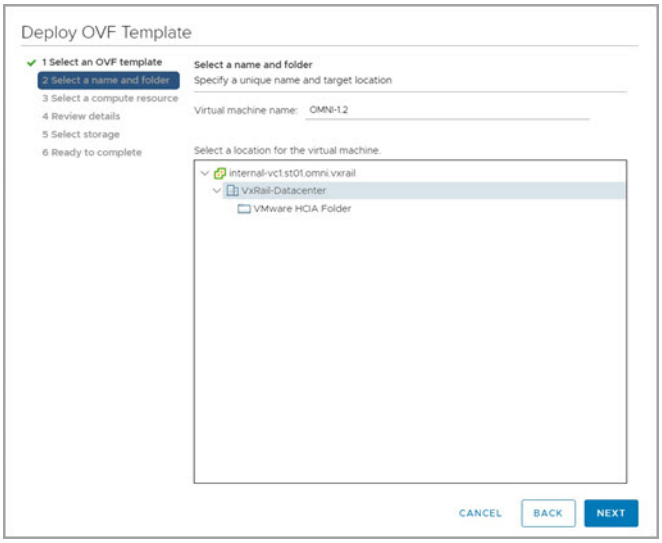

5. Select the destination compute resource, then click **Next**.

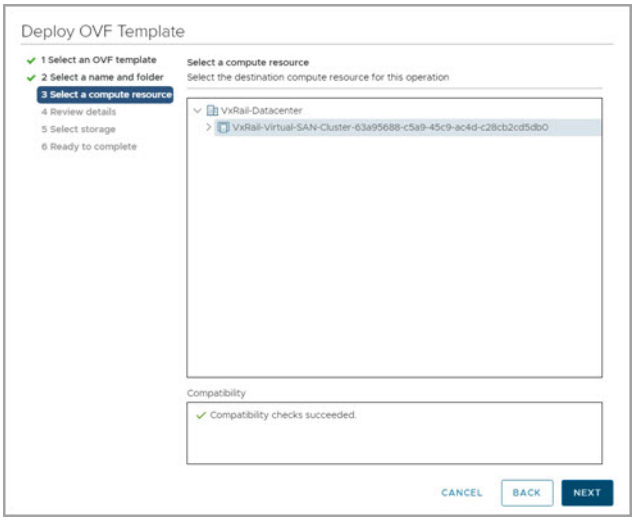

6. Review and verify the template details, then click **Next**.

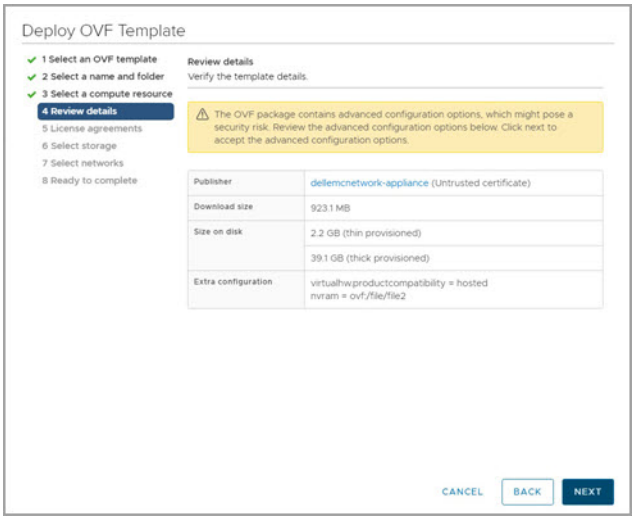

7. Accept the end-user license agreement (EULA), then click **Next**.

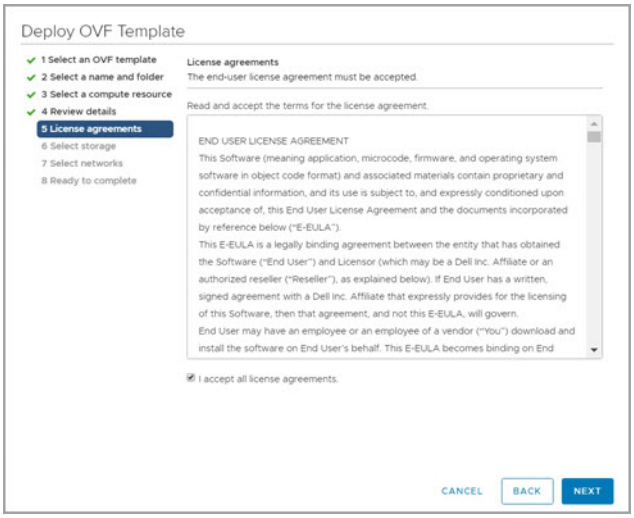

8. Select the VSAN datastore for the configuration and disk files, then click **Next**.

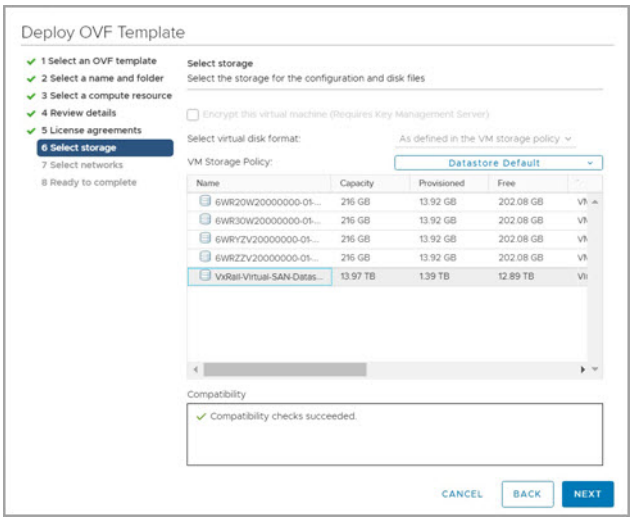

9. Select a destination network for each network source, then click **Next**. The VxRail Management Network must be assigned to the **VxRail internal Management network**. The default VLAN ID for this network is **3939**. The vCenter Server network must be connected to the port group where the vCenter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

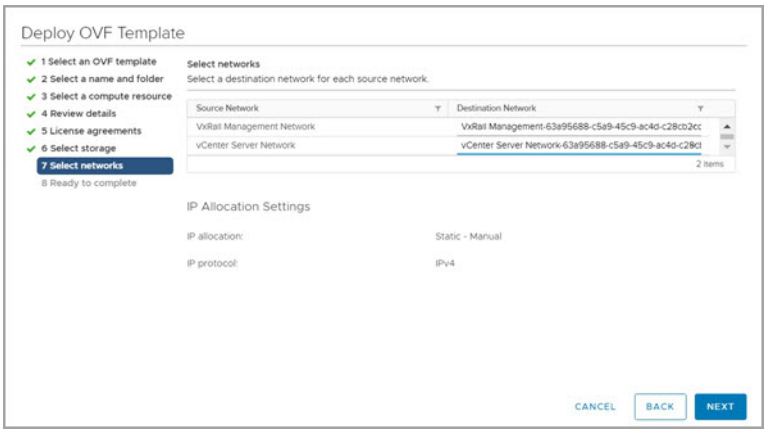

10. Click **Finish** to start creation of the VM.

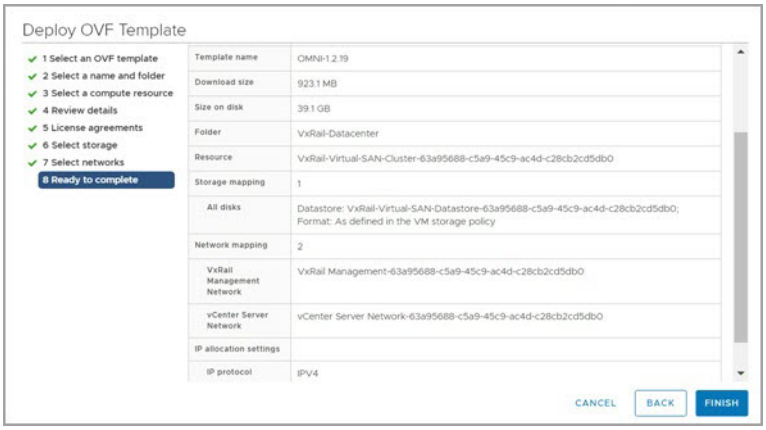

Power on OMNI VM

1. Scroll to the bottom of the window to view the status, and wait for the deployment to finish.

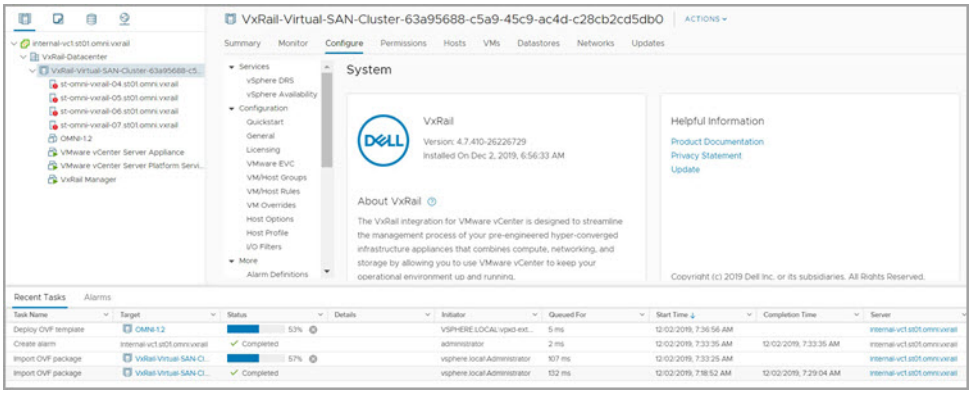

2. Select the **OMNI VM**, then select **Actions > Power On**.

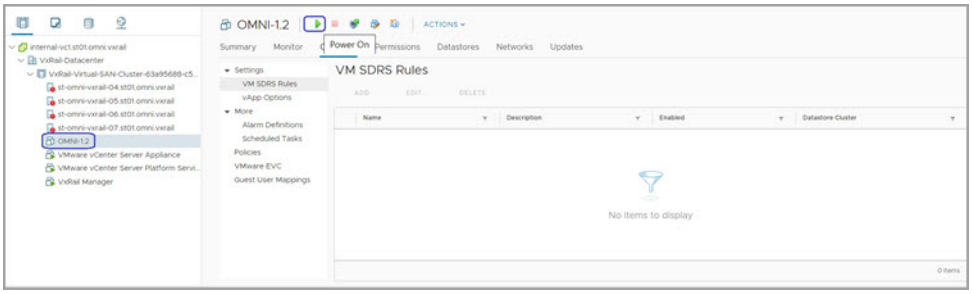

3. Select **Launch Web Console**.

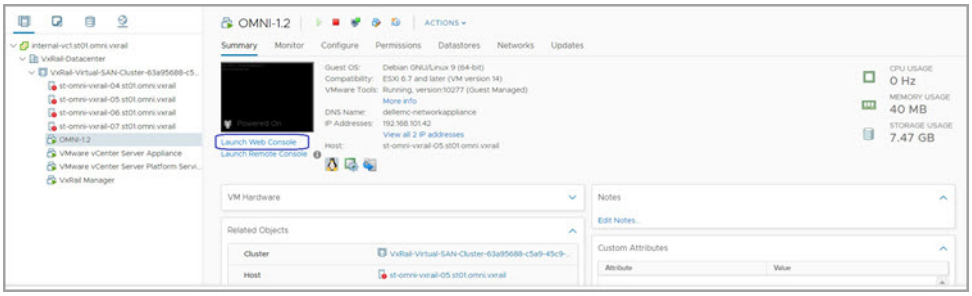

OMNI setup

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

Log in to VM console

You can configure OMNI through the VM console after you complete the authentication step. The VM console automatically closes after 10 minutes (default), and can be customized to meet your needs.

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

2. If this is a first-time login, change the password.

   After the passwords are successfully updated, self-signed certificates are created. You can change the certificates later with menu options.

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 SmartFabric menu. You can only access root user through the console.

Setup OMNI

This information describes how to set up the appliance with the required network interface configurations, and registration with vCenter and SmartFabric. A single OMNI VM instance supports up to 10 vCenters and 16 SmartFabric domains.

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

Network interface profile configuration

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

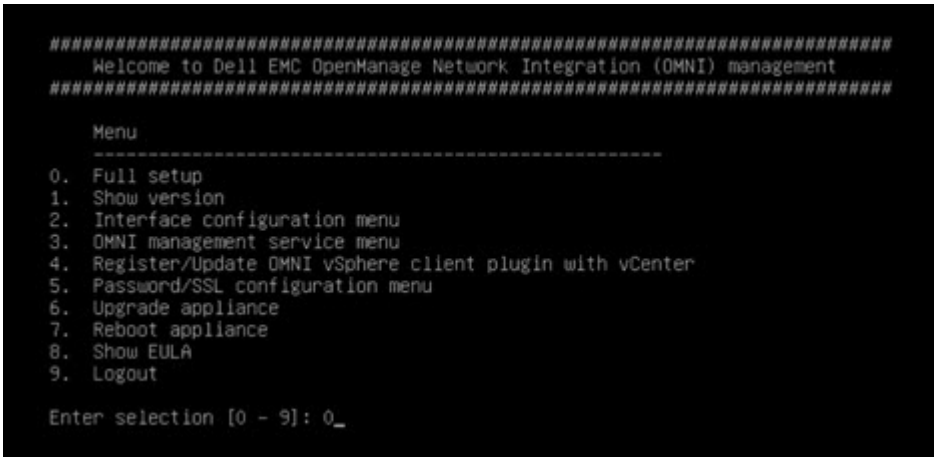

2. Select **Edit a connection**, then click **OK**.

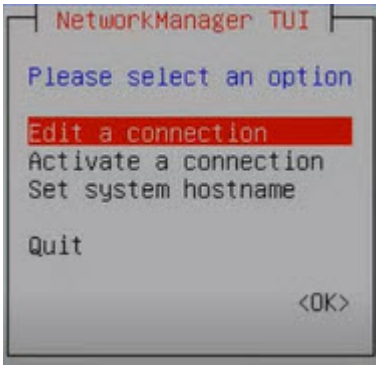

3. Select **Wired connection 1**, then click **Edit**.

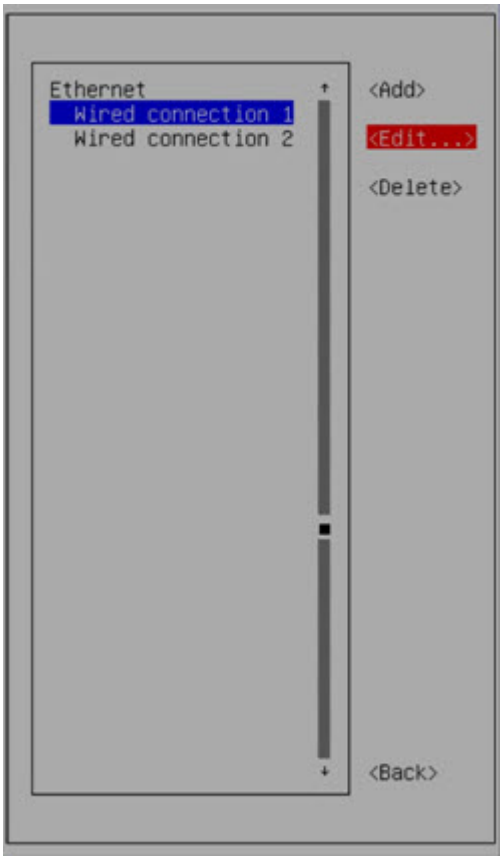

4. Verify Ethernet (ens160) is connected to the vCenter reachable network, then change the Profile name to **vCenter Server Network**.

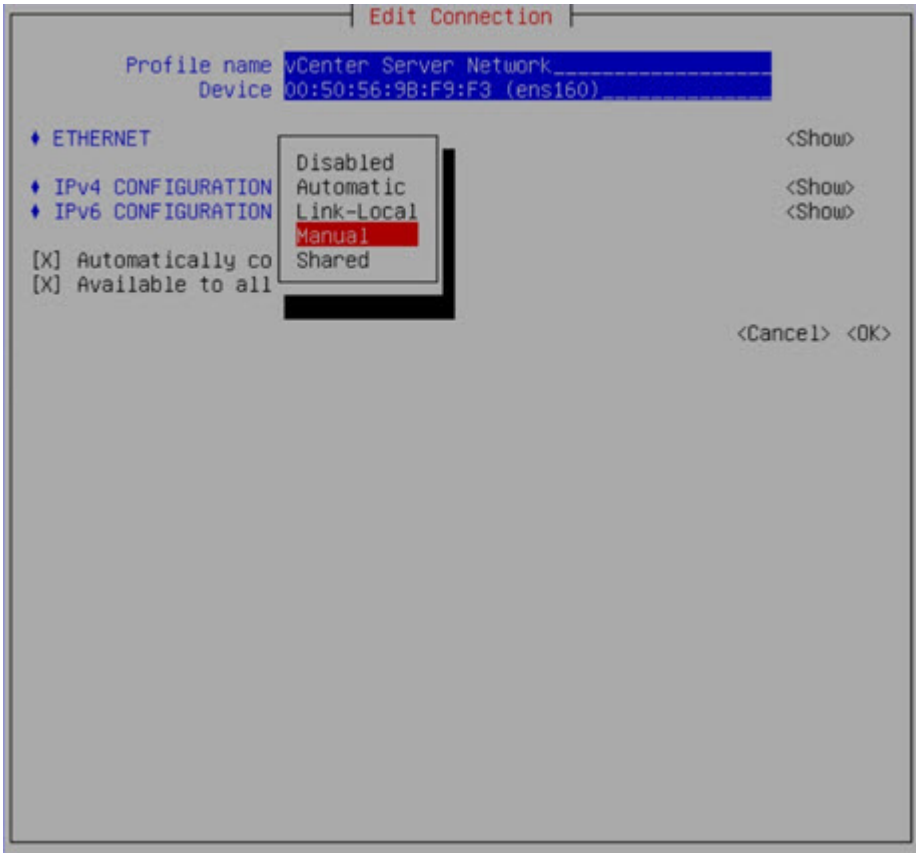

5. Change the IPv4 configuration from Automatic to Manual from the drop-down. You can choose Automatic or Manual IP address configuration.

   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.

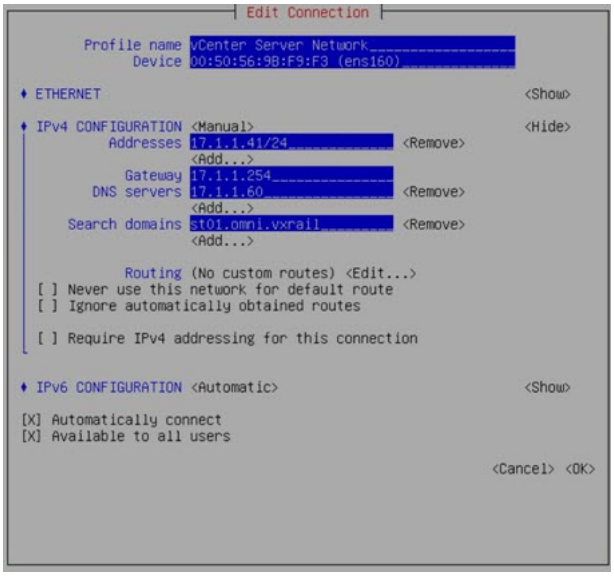

6. Click **Show** to the right of IPv4 configuration, then click **Add**.

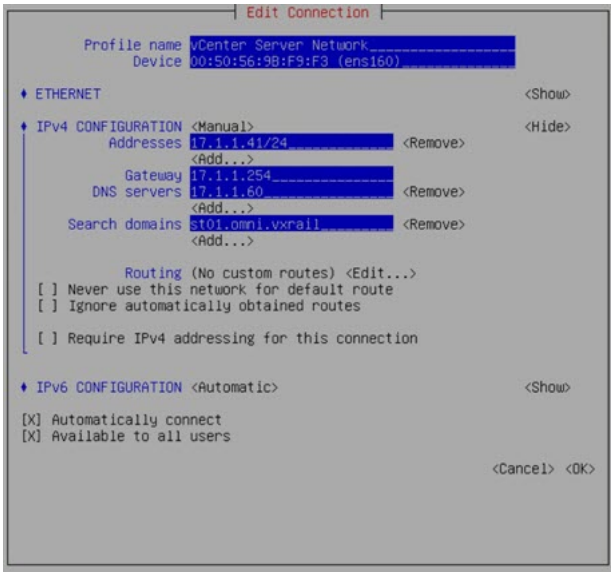

7. Set the Manual IPv4 address, Gateway address, DNS servers, Search Domains, then click **Edit** to the right of Routing.

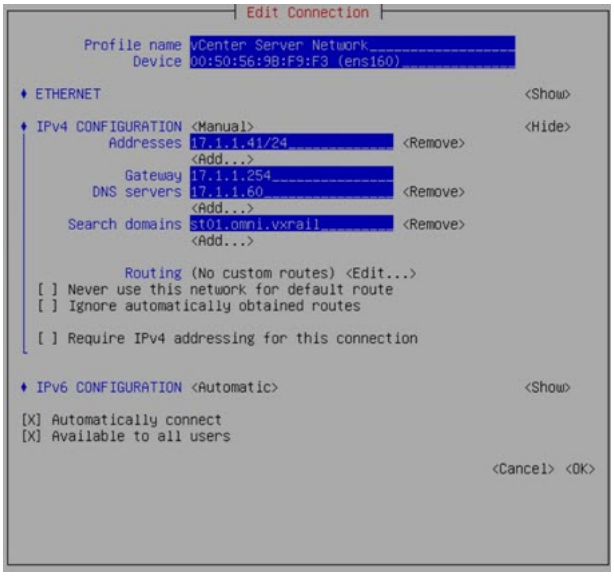

8. On IPv6 configuration, select **Ignore** for the IPv6 configuration, then click **OK**.

NOTE: IPv6 configuration is only required for an in-band network.

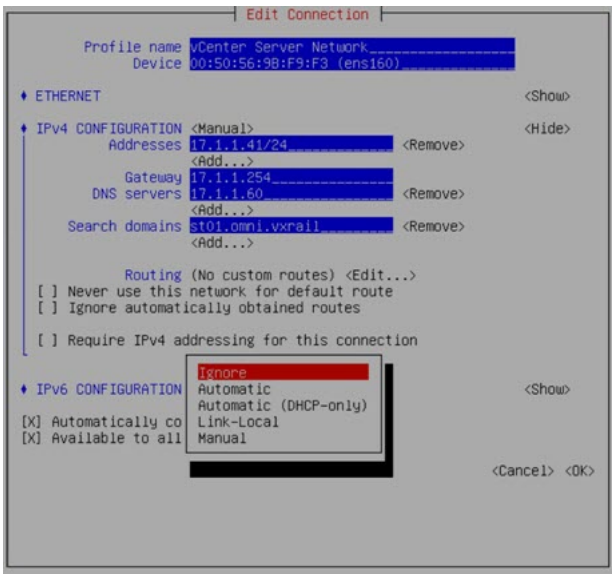

You are now ready to continue configuration.

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

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

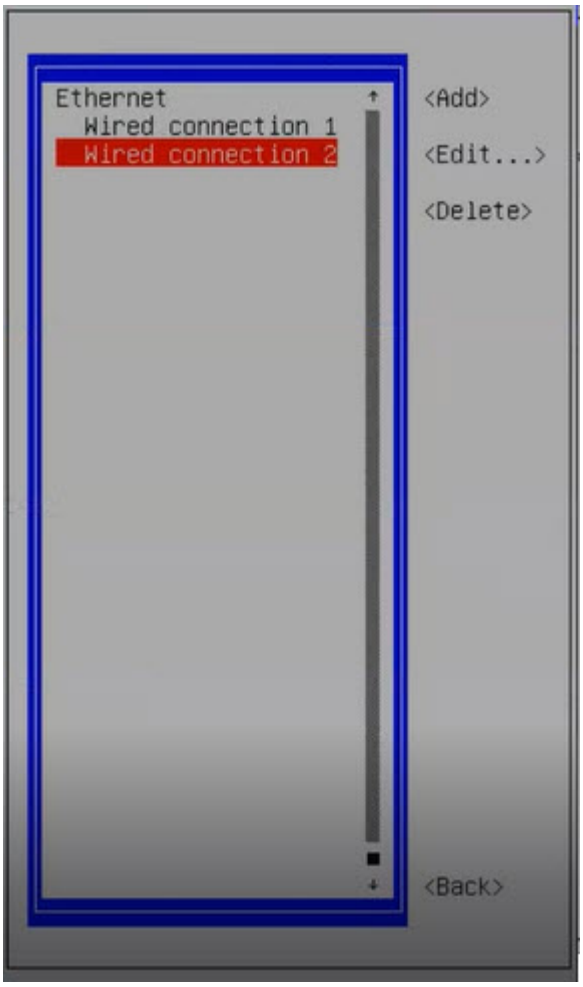

2. Rename Profile name to **VxRail Mgmt Network**.

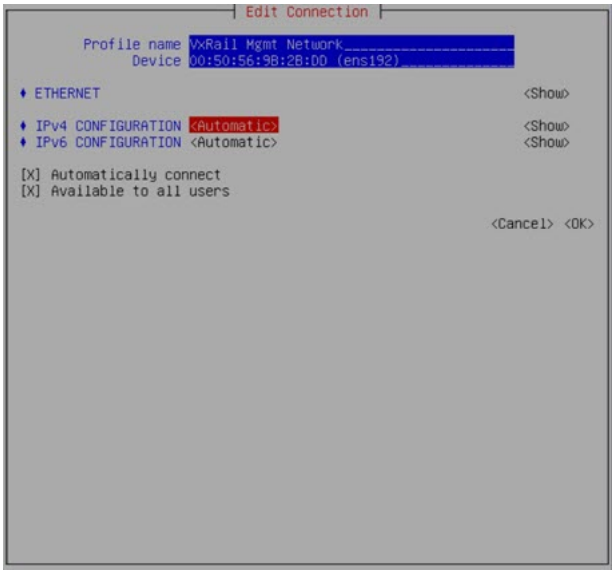

3. Select **Disabled** for the IPv4 configuration.

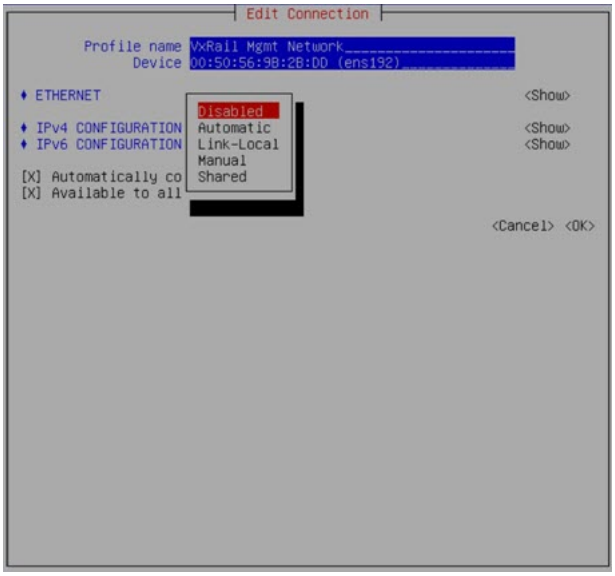

4. Select **Link-Local** for the IPv6 configuration.

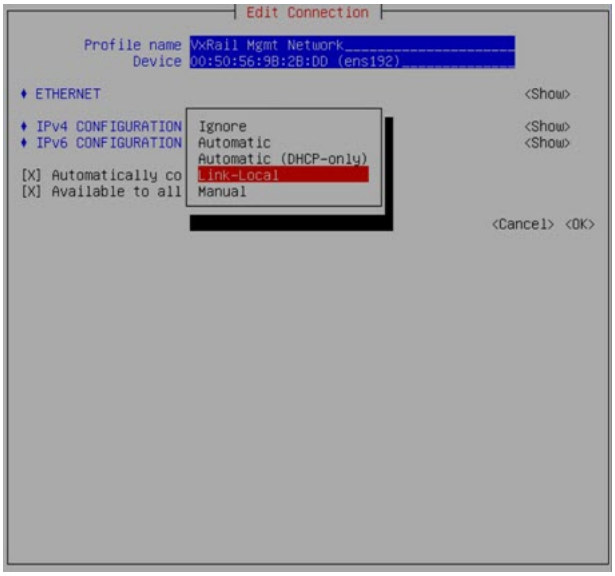

5. Click **Edit** to the right of Routing, then click **Add**.

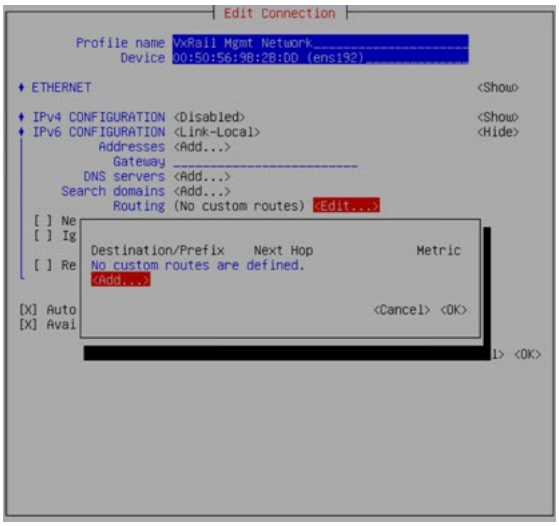

6. Enter the custom route as **fde1:53ba:e9a0:cccc::/64**, then click **OK**.

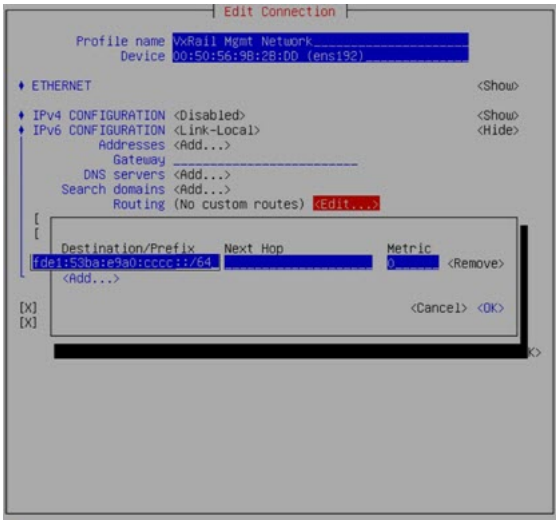

7. One custom route is now configured; click **OK**.

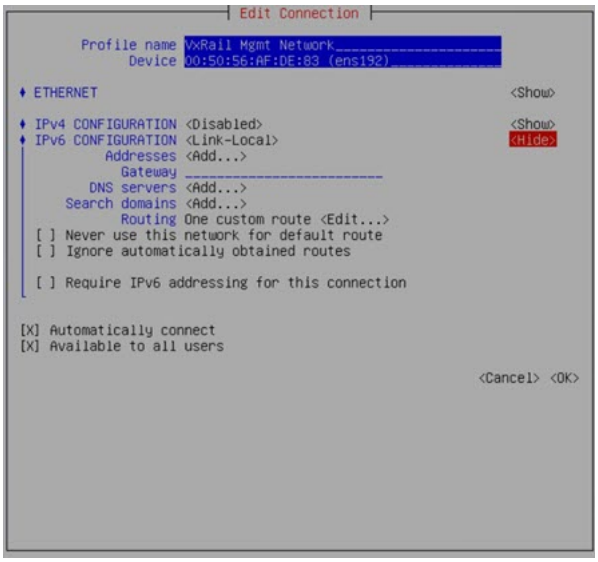

8. Click **Back** to activate the connection profiles.

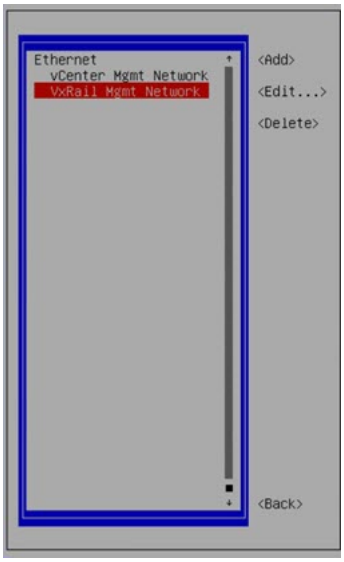

Activate connection profiles

NOTE: To populate DNS entries automatically, each profile must be deactivated, then activated.

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

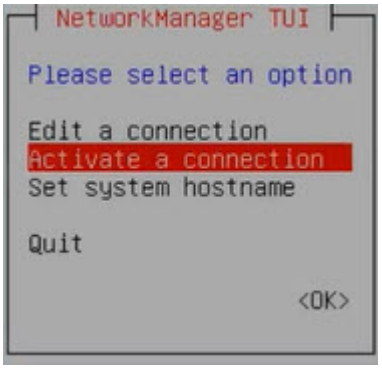

NOTE: If you make any changes while editing a connection, you must deactivate then activate the connection for the respective interface profile.

2. Select the **vCenter Server Network** profile, then click **Deactivate**; repeat for **VxRail Mgmt Network**.

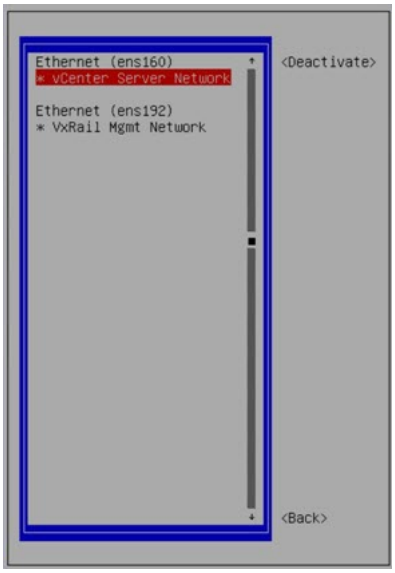

3. Select the **vCenter Server Network** profile, then click **Activate**; repeat for **VxRail Mgmt Network**.

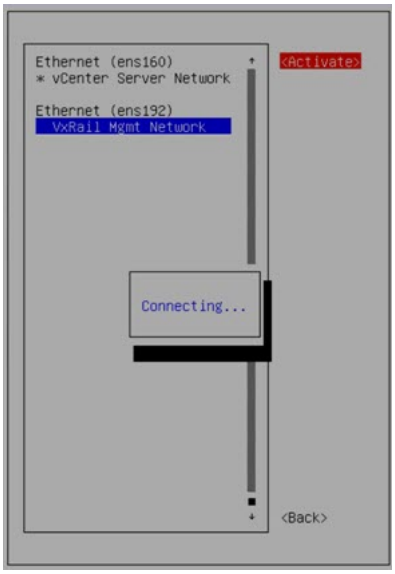

4. Click **Back**, then select **Set system hostname**, then click **OK**.

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, then click **OK**.

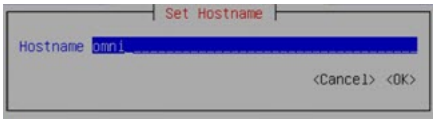

6. The hostname is now set; click **OK**.

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

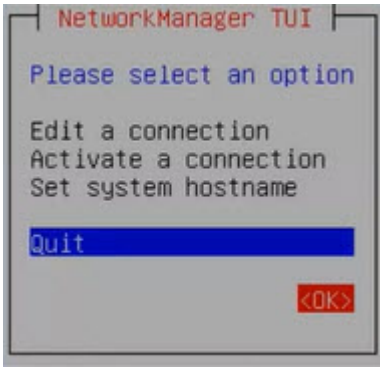

8. Enter a valid NTP Server IP address or hostname, then click **Enter**. Go to the next section for the SSL certificate installation.

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

OMNI vCenter client plug-in registration

This information describes how to register the vCenter plug-in, and SSL certificate management. SSL certificates have been automatically generated after the password is successfully updated. For more information, see Log into VM console.

NOTE: Multiple OMNI instances cannot be mapped to a single vCenter instance. If a situation where multiple VxRail clusters exist with their own respective fabric instances, it is recommended to map these fabric instances to a single vCenter instance. For example, VxRail cluster1 ideally has its own vCenter-1 VM instance, and the same is true for VxRail cluster 2 with its own vCenter-2 VM instance. In this case, OMNI-1 maps to vCenter-1, and OMNI-2 maps to vCenter-2.

If you do not want to create individual OMNI to vCenter mappings, you do have the option of mapping multiple fabric instances to a single OMNI mapped to a single or primary vCenter instance.

1. Select **N** to not install the SSL certificate now.

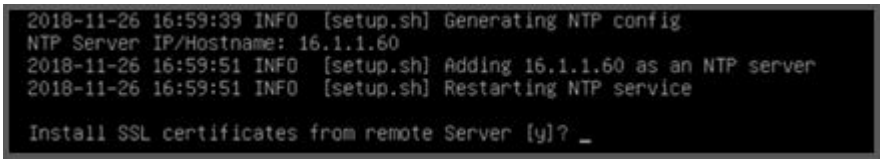

2. Select **4. Register/Update OMNI vSphere client plugin with vCenter**.

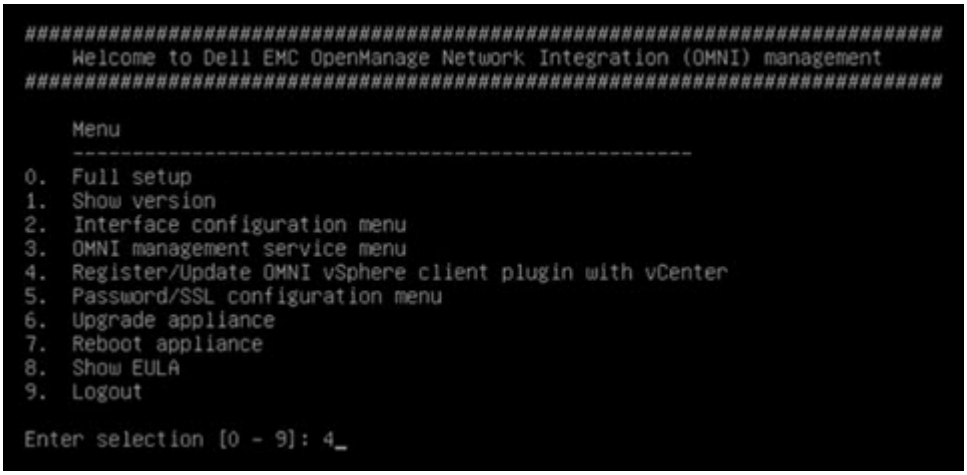

3. Enter the OMNI IP/FQDN for registration with the vCenter instance.

NOTE: We recommend using FQDN instead of the IP address of OMNI.

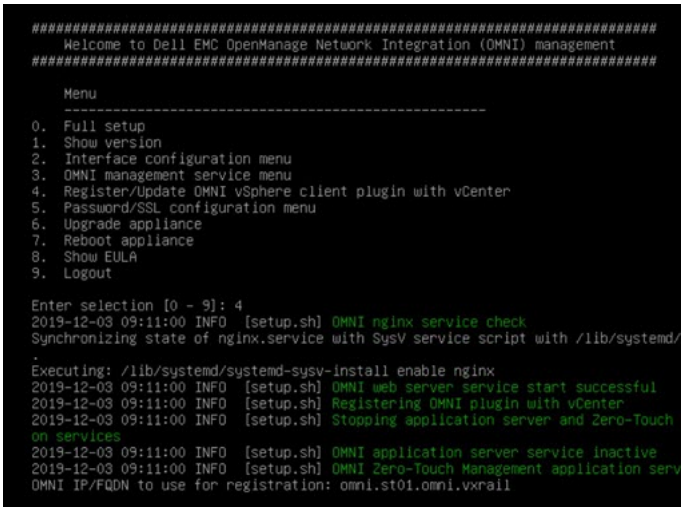

4. Enter the vCenter Server FQDN, vCenter Server username and vCenter Server password. Repeat this step to register each vCenter instance (up to 10).

```
########################################################################
    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.  Register/Update OMNI vSphere client plugin with vCenter
5.  Password/SSL configuration menu
6.  Upgrade appliance
7.  Reboot appliance
8.  Show EULA
9.  Logout

Enter selection [0 - 9]: 4
2019-12-03 09:11:00 INFO  [setup.sh] OMNI nginx service check
Synchronizing state of nginx.service with SysV service script with /lib/systemd/systemd-sysv-install
.
Executing: /lib/systemd/systemd-sysv-install enable nginx
2019-12-03 09:11:00 INFO  [setup.sh] OMNI web server service start successful
2019-12-03 09:11:00 INFO  [setup.sh] Registering OMNI plugin with vCenter
2019-12-03 09:11:00 INFO  [setup.sh] Stopping application server and Zero-Touch Management applicati
on services
2019-12-03 09:11:00 INFO  [setup.sh] OMNI application server service inactive
2019-12-03 09:11:00 INFO  [setup.sh] OMNI Zero-Touch Management application service inactive
OMNI IP/FQDN to use for registration: omni.st01.omni.vxrail
Appliance IP : omni.st01.omni.vxrail
vCenter server FQDN: internal-vc1.st01.omni.vxrail
vCenter server username: administrator@vsphere.local
vCenter server password:
```

5. The OMNI application server services start successful; press **[enter]** to continue.

```
3.  OMNI management service menu
4.  Register/Update OMNI vSphere client plugin with vCenter
5.  Password/SSL configuration menu
6.  Upgrade appliance
7.  Reboot appliance
8.  Show EULA
9.  Logout

Enter selection [0 - 9]: 4
2019-12-03 09:11:00 INFO  [setup.sh] OMNI nginx service check
Synchronizing state of nginx.service with SysV service script with /lib/systemd/systemd-sysv-install
.
Executing: /lib/systemd/systemd-sysv-install enable nginx
2019-12-03 09:11:00 INFO  [setup.sh] OMNI web server service start successful
2019-12-03 09:11:00 INFO  [setup.sh] Registering OMNI plugin with vCenter
2019-12-03 09:11:00 INFO  [setup.sh] Stopping application server and Zero-Touch Management applicati
on services
2019-12-03 09:11:00 INFO  [setup.sh] OMNI application server service inactive
2019-12-03 09:11:00 INFO  [setup.sh] OMNI Zero-Touch Management application service inactive
OMNI IP/FQDN to use for registration: omni.st01.omni.vxrail
Appliance IP : omni.st01.omni.vxrail
vCenter server FQDN: internal-vc1.st01.omni.vxrail
vCenter server username: administrator@vsphere.local
vCenter server password:
2019-12-03 09:36:12,422 Extension registration succeed with: internal-vc1.st01.omni.vxrail
2019-12-03 09:36:12 INFO  [setup.sh] Starting application server and Zero-Touch Management applicati
on services
Created symlink /etc/systemd/system/multi-user.target.wants/dellvcenterapp.service → /etc/systemd/sy
stem/dellvcenterapp.service.
2019-12-03 09:36:12 INFO  [setup.sh] OMNI Zero-Touch Management application service start successful
Created symlink /etc/systemd/system/multi-user.target.wants/vcenterappgunicorn.service → /etc/system
d/system/vcenterappgunicorn.service.
2019-12-03 09:36:12 INFO  [setup.sh] OMNI application server service start successful (plugin versio
n 1.2.19).
press [enter] to continue...
```

6. Select **9. Logout**.

```
########################################################################
    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.  Register/Update OMNI vSphere client plugin with vCenter
5.  Password/SSL configuration menu
6.  Upgrade appliance
7.  Reboot appliance
8.  Show EULA
9.  Logout

Enter selection [0 - 9]: 9_
```

NOTE: You cannot register the same vCenter instance from another OMNI plug-in.

Unregister vCenter with OMNI

1. Select **vCenter Server > Menu > OpenManage Network Integration > Service Instance**, then select the specific **Service Instance > Summary > vCenter Credentials**.

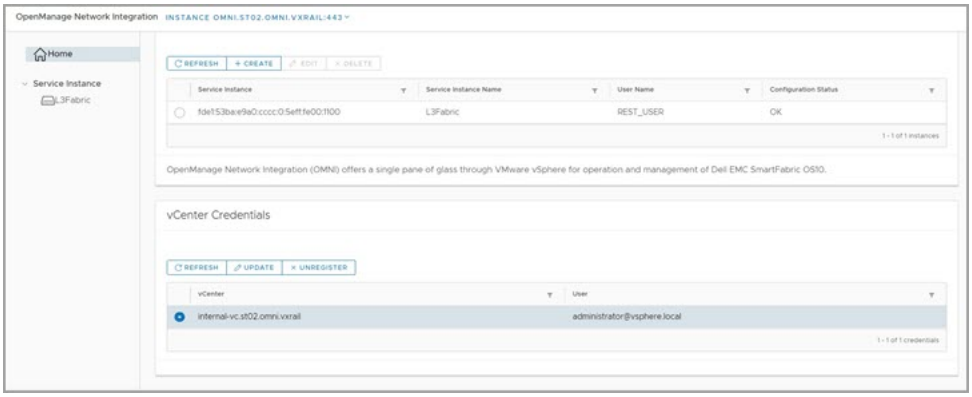

2. Click **Unregister**.

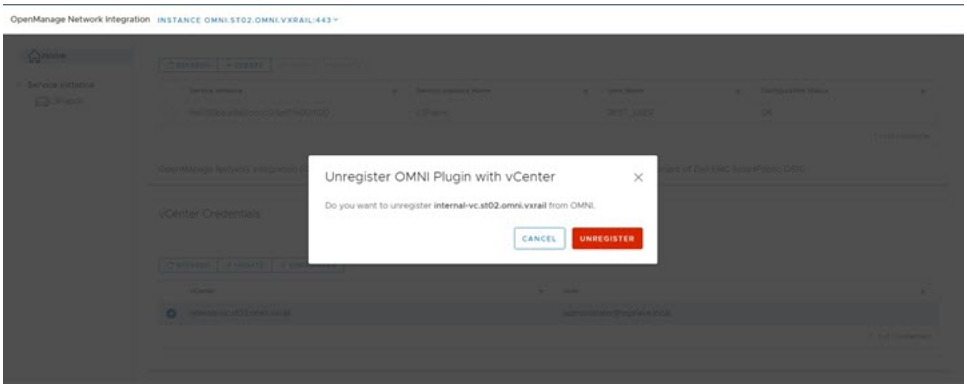

vCenter credential update of registered vCenter

1. Click **Update** to update the credentials of the existing vCenter.

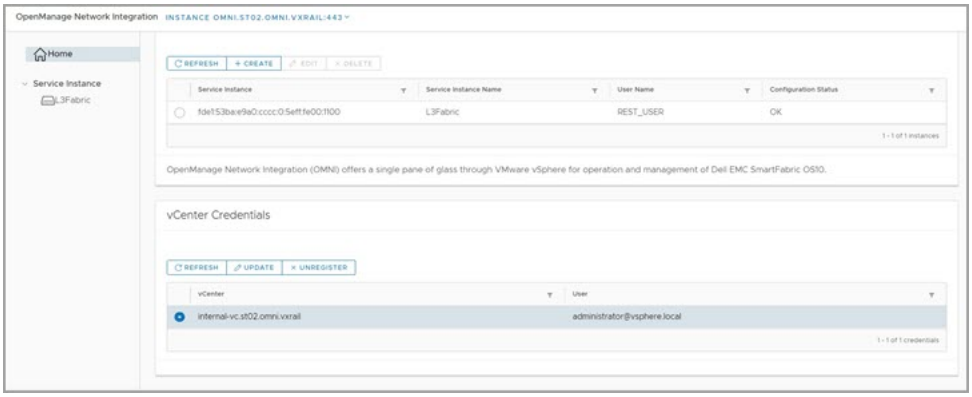

2. Enter the required information to edit (user and vCenter fields are automatically populated), then click **Update**.

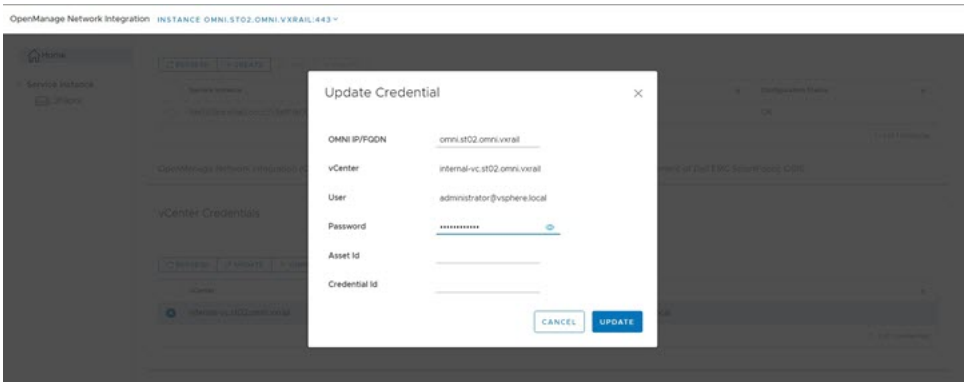

3. Update success; close the message.

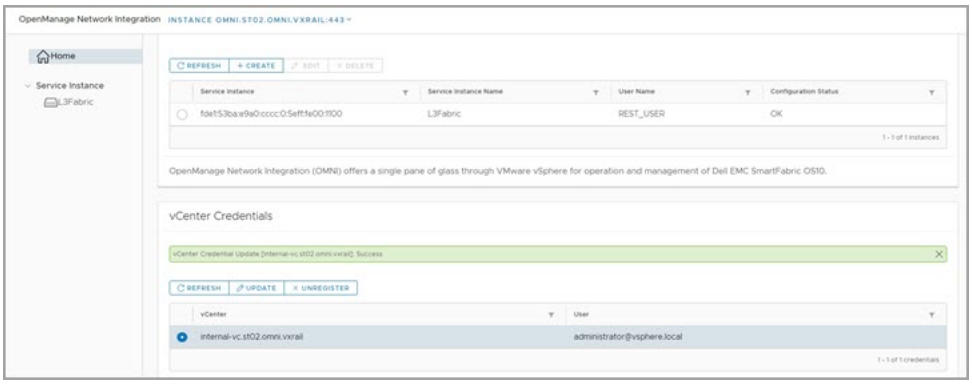

4

OMNI virtual appliance management menu

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

Table 3. OMNI virtual appliance management menu

Menu option	Submenu option	Description
1. Show version		Display OMNI virtual appliance and plug-in version
2. Interface configuration menu	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/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 server, application server, and VLAN automation process
	2. View OMNI management service	Display status of OMNI provided services
	3. Stop OMNI management service	Stop OMNI web server, application server, and VLAN automation process
	4. Restart OMNI management service	Restart OMNI web server, application server, and VLAN automation process
	5. Create support bundle	Create OMNI support bundle archive and save to download location. **It is recommended that you use the OMNI user interface support page to generate and download the support bundle**.
	6. Change application log level	Display current log-levels, and configure DEBUG or ERROR log-levels. **It is recommended that you use the OMNI user interface support page to change the OMNI log-levels**.
	7. Exit	

Menu option	Submenu option	Description
4. Register/update OMNI vSphere client plug-in with vCenter		Register OMNI with vCenter; enter OMNI IP/hostname, vCenter IP/hostname, vCenter administrator user (administrator @vsphere.local), and vCenter password. **It is recommended that you use hostnames**.
5. Password/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**.
	3. Generate self-signed SSL certificates	Replace existing OMNI appliance self-sign certificate. **After SSL certificate installation completes, you need to re-register OMNI with the vCenter**.
	4. Install SSL certificates from remote server	Replace OMNI certificates with the certificate that is located on the remote server using SCP/FTP. **After SSL certificate installation completes, you need to re-register OMNI with the vCenter**
	5. Exit	
6. Upgrade appliance		Upgrade the OMNI appliance. **Upgrade appliance can only be used for same version OMNI patch installation**.
7. Reboot appliance		Reboot the OMNI appliance
8. Show EULA		Display the OMNI end user license agreement (EULA)
9. Logout		Log out as the `admin` user

5

Access to OMNI portal

This information describes how to access SmartFabric vCenter through the vSphere Client. A shortcut is available from the vSphere Client left-pane within the menu dropdown and shortcuts view.

Access OMNI portal using registered vCenter

NOTE: Before you use the plug-in, you must set up a Dell EMC SmartFabric appliance in vSphere. Once you register OMNI with vCenter, an OMNI plug-in is available in the vCenter.

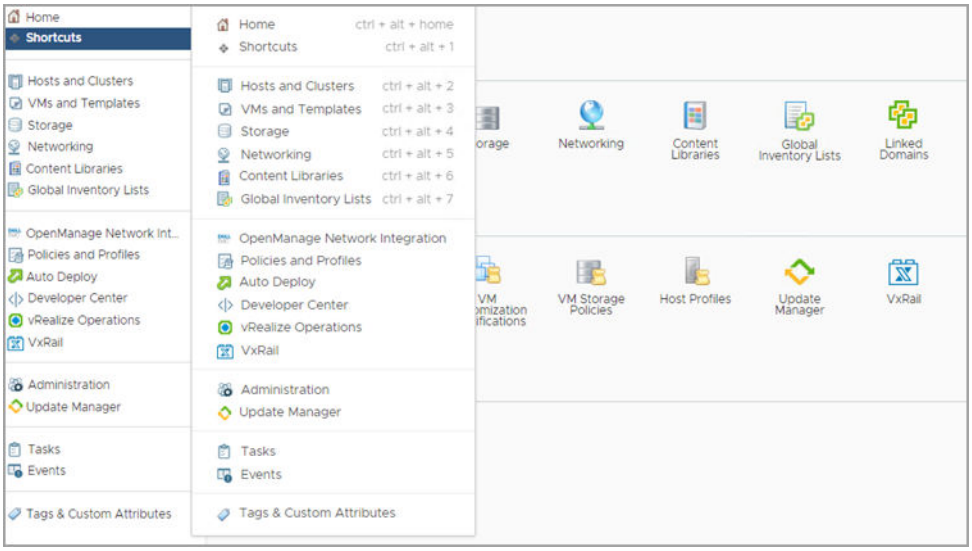

When you select SmartFabric, the home page displays information about the SmartFabric domains being managed. This page also allows you to update extensions if available. Information includes:

- Service instance
- vCenter credentials

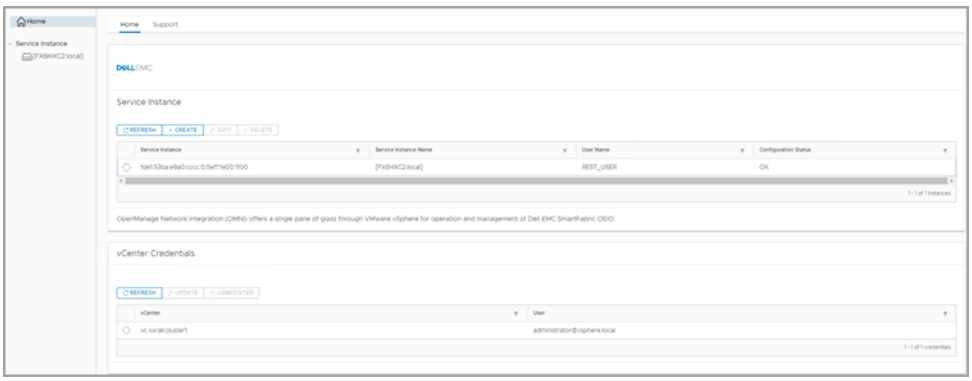

OMNI portal access using OMNI appliance IP

NOTE: If you are using a stand-alone generic ESXi host deployment, this information does not apply and you can skip this section.

Once the vCenter registration steps are complete, you can log in to the OMNI UI using the OMNI stand-alone page. This information describes how to access the OMNI UI from a browser.

1. Open a browser session, go to **https://*OMNI_IP*/delawareos10** with the IP configured during setup.
2. Enter the **username** and **password** for the OMNI VM, then click **Sign In**.

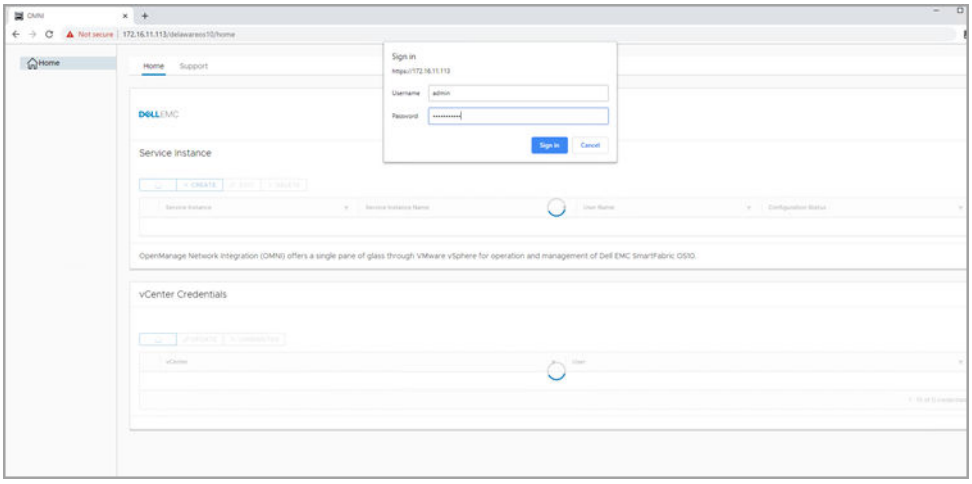

Once the username and password are authenticated, the OMNI page displays.

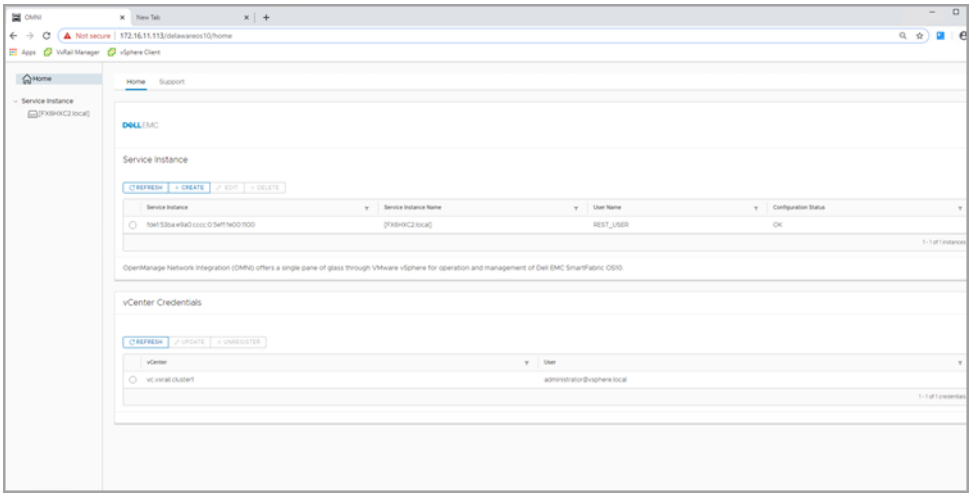

6

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.

NOTE: You cannot delete `PortGroups` on a VSS/DVS, or delete the VSS/DSS entirely as it clears all unused networks from the SmartFabric `ServerInterfaces`.

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.

We recommend keeping the DVS uplink in Trunking mode and configure 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.

NOTE: When any server port is removed from an uplink on DVS/vDS, the VLAN association details stored on the switches are not deleted. It is recommended to remove the network association configuration manually from the server interface profile.

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 4. vCenter `PortGroup` VLAN automation of identified ESXi host

vCenter action	SmartFabric action by OMNI
Add/update `PortGroup`: VLAN of VSS/DVS	• Create network of `PortGroup` VLAN • Add network to SmartFabric `ServerInterface`
Remove `PortGroup` from VSS/DVS	Remove unused networks from SmartFabric `ServerInterface`

NOTE: OMNI automation is not designed to delete unused `ServerInterfaces` of SmartFabric.

SmartFabric networks consolidation by OMNI

1. Collect all networks of registered SmartFabric.
2. Collect networks of `ServerInterface` of registered SmartFabric.
3. Identify SmartFabric networks created by the OMNI user interface, and SmartFabric networks that are not created by the OMNI user interface.
4. Append networks that are not created by the OMNI user interface to SmartFabric uplink of 'Default'/'CreateOnly' type
5. Find unused networks; SmartFabric networks not created by the OMNI user interface, and not used by the SmartFabric `ServerInterface` and SmartFabric uplinks.

6. Delete unused networks from the SmartFabric.

NOTE: A 'Default'/'CreateOnly' uplink can be configured on the SmartFabric through the OMNI Uplink configuration page. For more information, see OMNI SmartFabric management.

7

OMNI SmartFabric management

This information provides details of how OMNI helps you manage SmartFabric OS10 with OMNI. Also explained is how you can add and configure SmartFabric instances that you want to manage using OMNI.

Add SmartFabric instance

This information describes how to add SmartFabric instances in OMNI.

1. Go to the OMNI portal.
2. Locate the Master IP of the SmartFabric cluster by logging into OS10 SmartFabric, then enable the switch.

```
OS10# show smartfabric cluster

-------------------------------------------------------
CLUSTER DOMAIN ID : 100
VIP               : fde2:53ba:e9a0:cccc:0:5eff:fe00:1100
ROLE              : BACKUP
SERVICE-TAG       : FX6HXC2
MASTER-IPV4       : x.x.x.x
PREFERRED-MASTER  : false
-------------------------------------------------------
```

3. Click **Create** to manually add the Master IP of the SmartFabric Service instance.

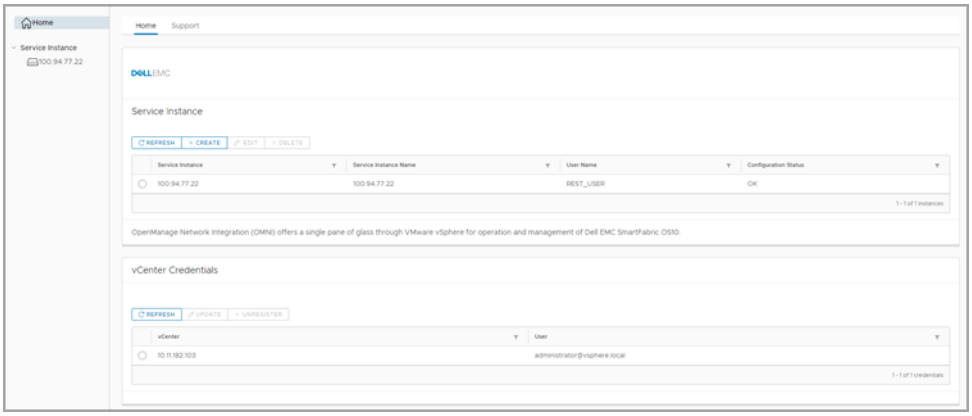

4. Enter the service instance information, then click **Add**.

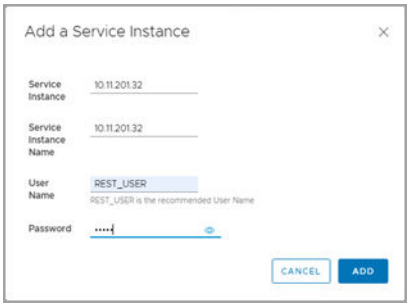

5. Service instance creation success; close the message.

Configure 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), find the SmartFabric IPv6 VIP autodiscovered by OMNI. For complete information about discovery, see *mDNS service* in SmartFabric Services.

You must edit this autodiscovered SmartFabric instance for the REST_USER password to complete the configuration.

1. Go to the OMNI portal.
2. Select **Auto Discovered VIP**, then click **Edit**.

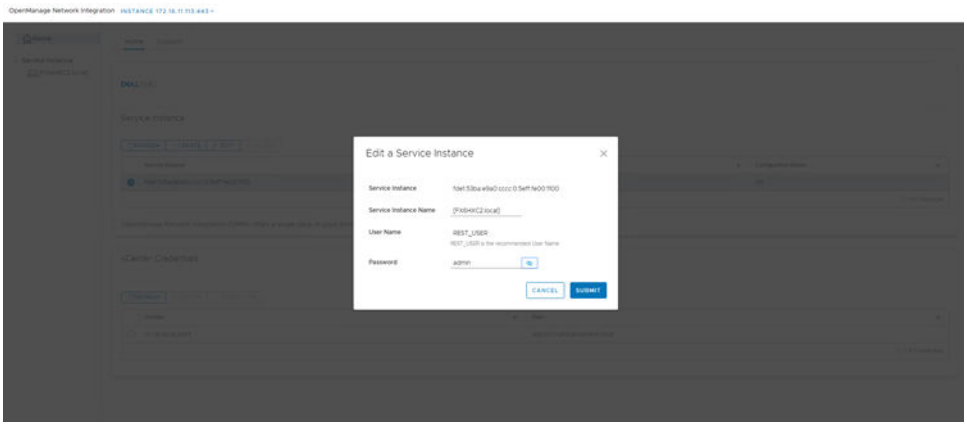

NOTE: During VxRail initial deployment, the system forces you to change the password. If you forget the REST_USER password, contact Dell support to reset REST_USER password.

3. Enter the service instance information, then click **Submit**.

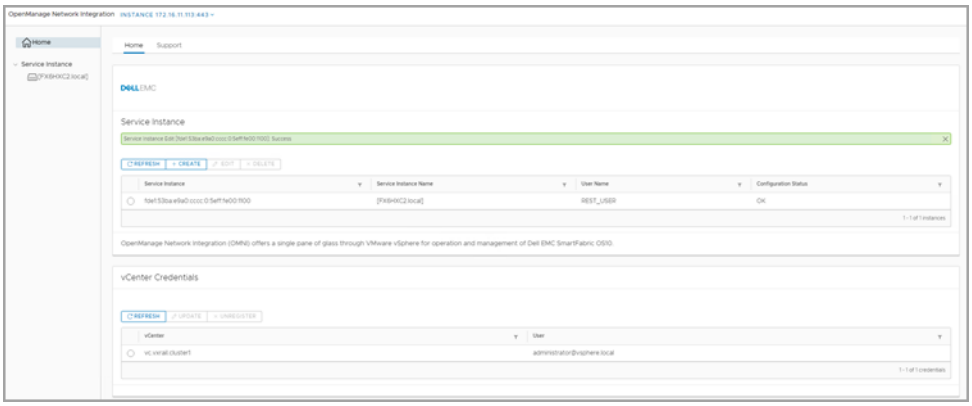

4. Service instance configuration success; close the message.

Topics:

- Summary
- Topology
- Switches
- Server interface
- Uplink
- Network
- Host network inventory

Summary

This information describes the selected fabric summary including fabric nodes, racks, and actions.

Select the Service Instance to view details on each fabric. Click **Domain** at any time to update the fabric details.

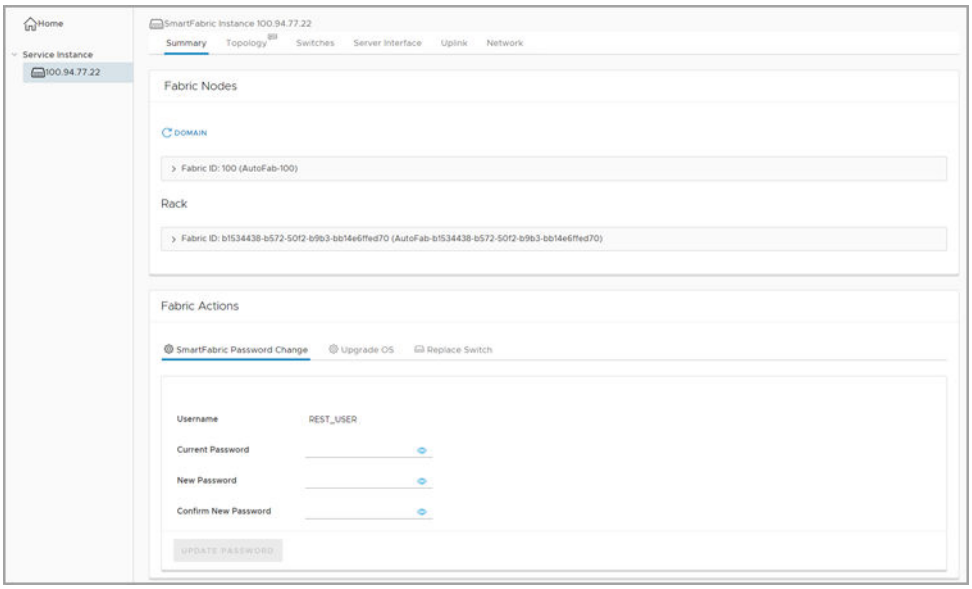

Fabric nodes

This information describes the selected Spine structure which is arranged in racks. Select the Fabric ID to view the fabric details. Each Spine fabric has corresponding switches. Each switch includes status (online or offline), name, model, version, role, and IP. Click **Domain** at any time to update the fabric details.

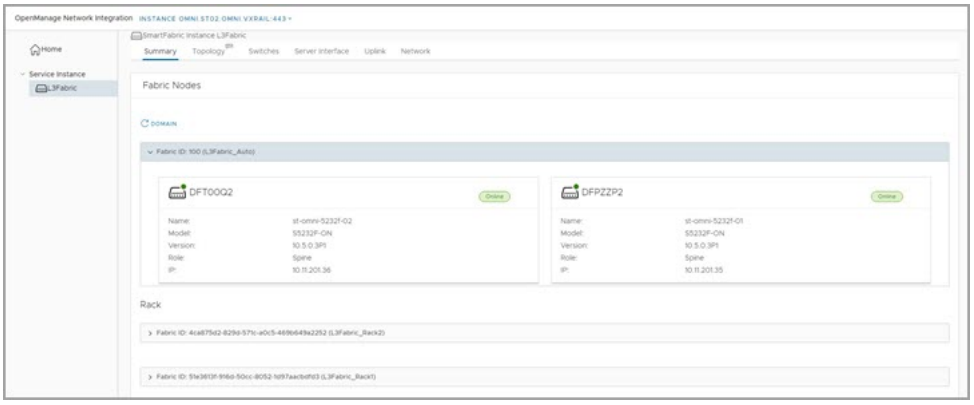

Rack

This information describes the rack which contains logical groupings of switches. For each rack, there is a fabric, and another fabric which contains all the information for the leafs. All leaf nodes are associated with one fabric. If the rack contains three fabrics, it contains two rack fabrics (one for each rack), and one fabric that contains all the information about the leaf.

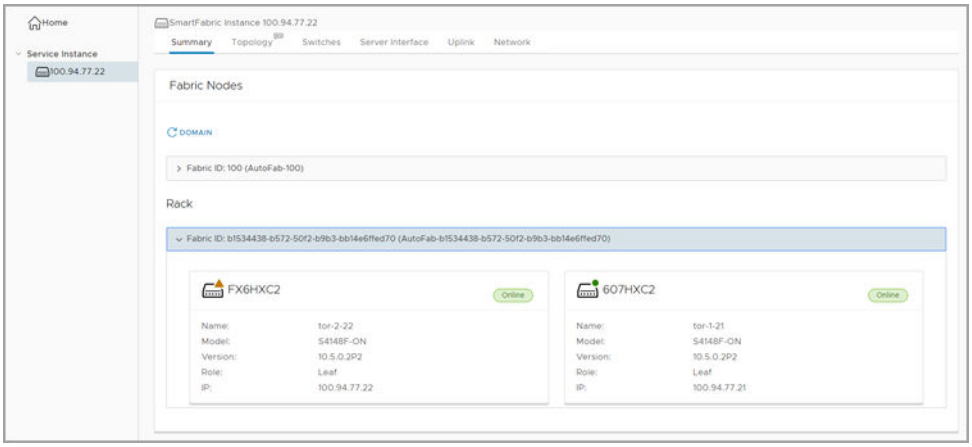

Fabric actions

You can do the following from the fabric actions:

- Change SmartFabric password.
- Upload SmartFabric OS10 to OMNI VM.
- Delete SmartFabric OS10 image.

  For more information about upload and delete SmartFabric OS10 image, see Upgrade SmartFabric OS.
- Replace switch on a network fabric.

  For complete information about replacing a switch, see Replace switch on network fabric.

SmartFabric password change

1. Enter the current password for the REST_USER, the new password, confirm the new password, then click **Update Password**.

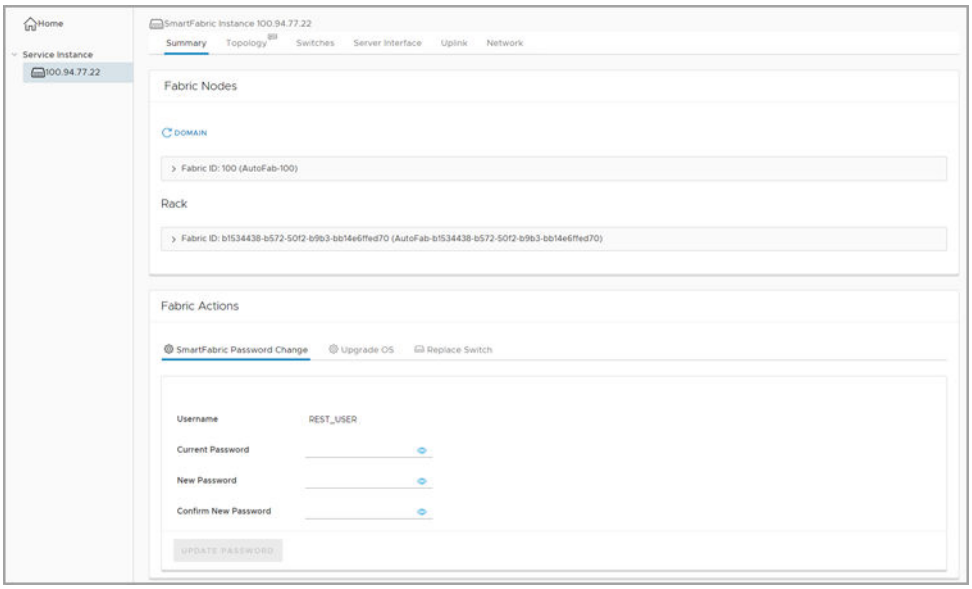

2. Password update success; failure tooltip notification message displays.

Topology

This information explains how to display the fabric topology and switch information. Switch information includes destination interface, destination switch, interface ID, interface name, switch ID, and type.

NOTE: This feature is marked beta for this release.

1. Select the **Service Instance**, then select **Topology**.

2. Mouse over a fabric to display switch details.

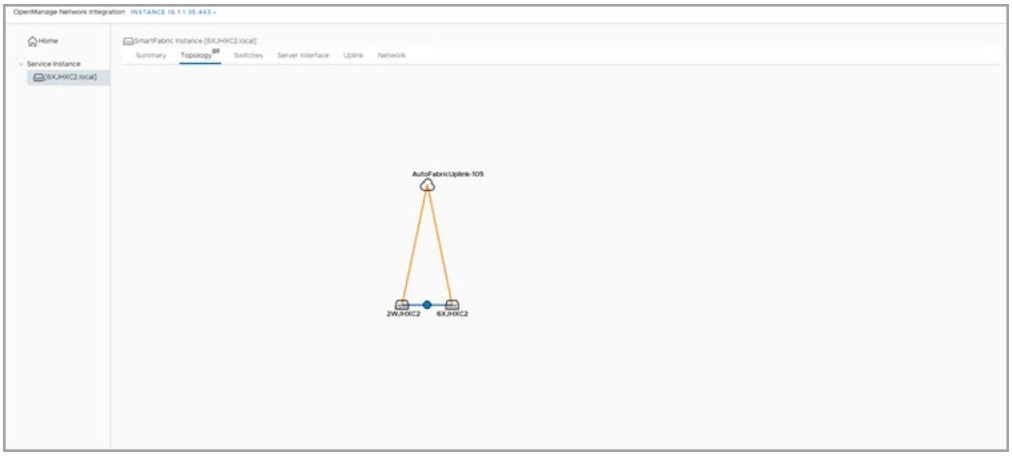

3. Select any switch to view the network information.

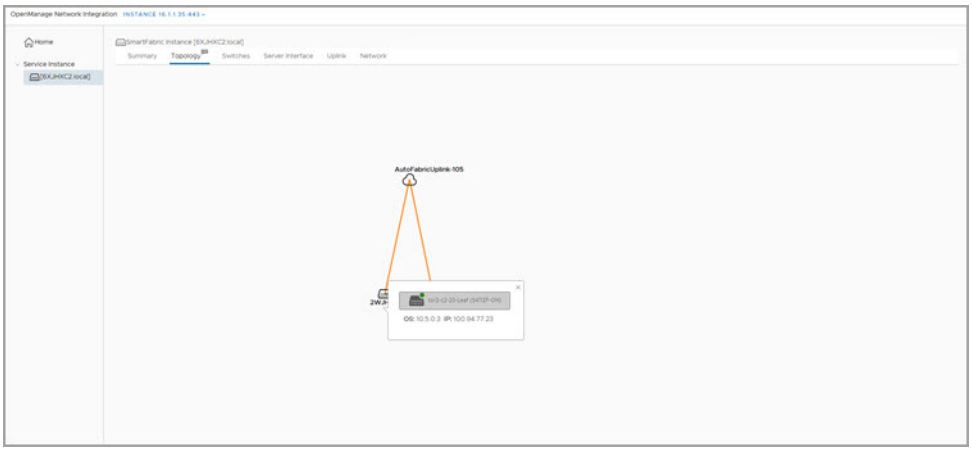

Switches

This information describes how to display fabric for all Spine and Leaf switches.

Switch information includes destination interface and breakout profile. Select a switch to view details about the specific switch including breakout details and corresponding information about interface status, MTU, and type.

1. Select the **Service Instance**, then select **Switches**.
2. Select the arrow to view the spine switch properties.

3. Select the switch.

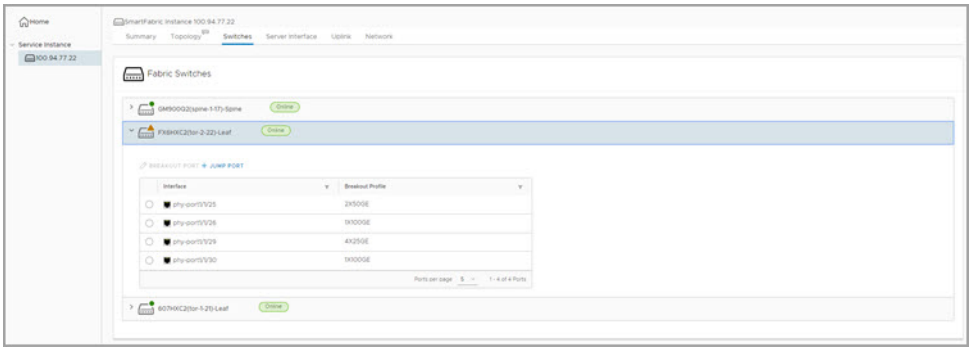

4. Use the arrow to view the properties.

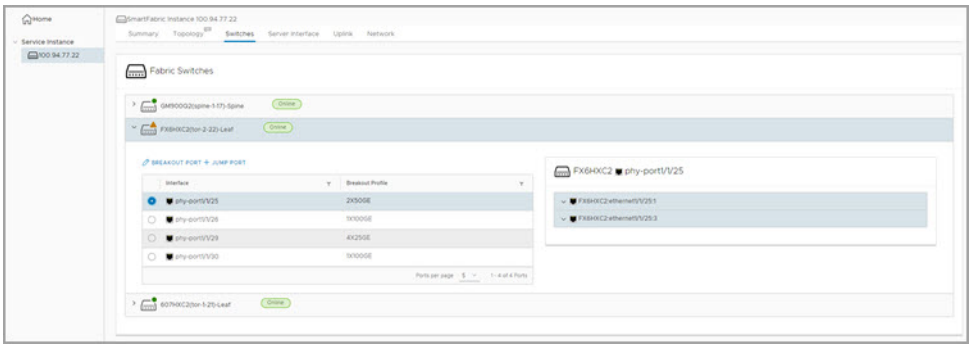

5. Select the interface/port-group to view properties on the right.

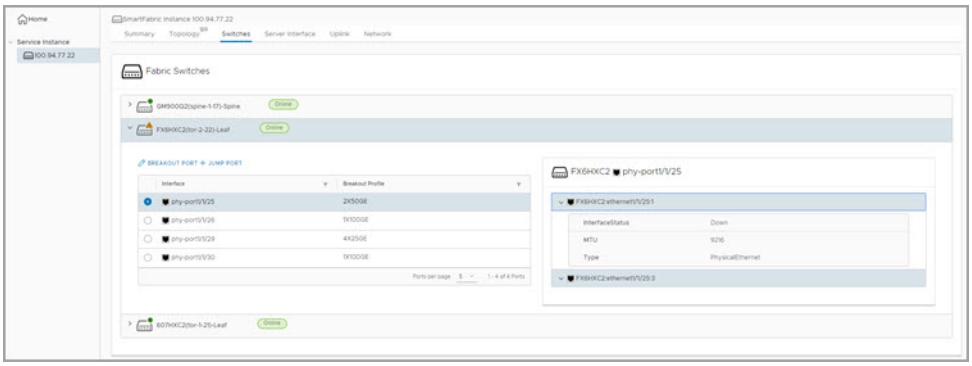

Breakout port

NOTE: The auto breakout feature is enabled by default in Spine switches. OMNI will not provide an option to break out Spine switches.

1. Select **Breakout port**, select the **Breakout Mode** from the drop-down, then click **Submit**.

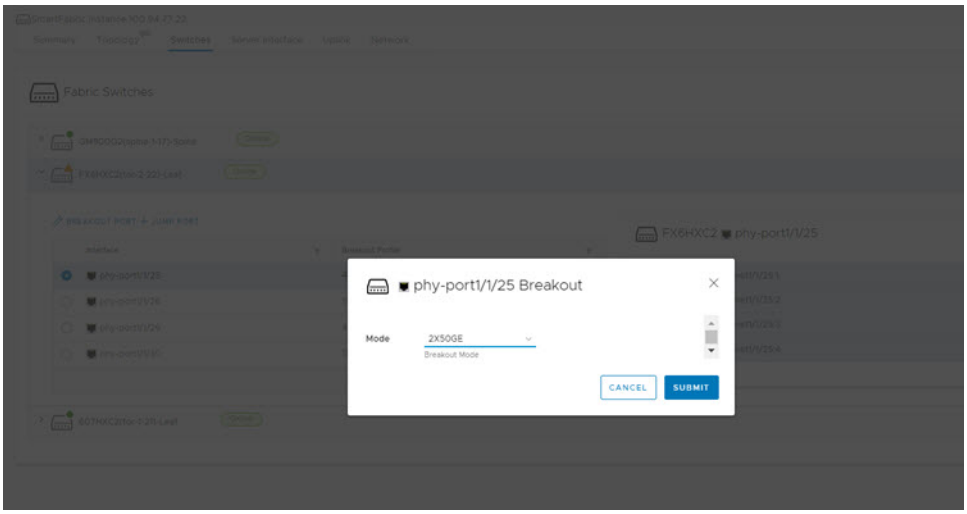

2. Breakout port successfully configured; success/failure tooltip notification message displays.
3. Select the number of ports to view by page.

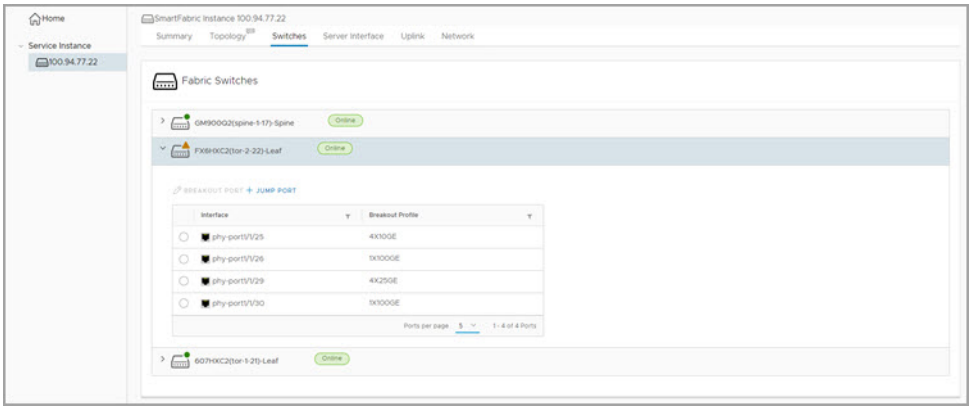

Add jump port

1. Select the switch to view the Leaf switch properties and the current jump port, then click **Jump Port**.

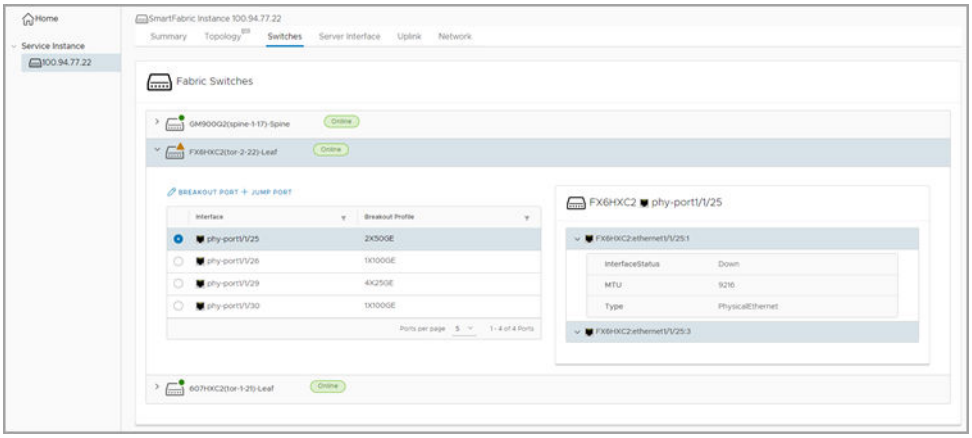

2. Enter the **Name** of the new jump port, select the **Interface Name**, select the **Untagged Network**, then click **Add**.

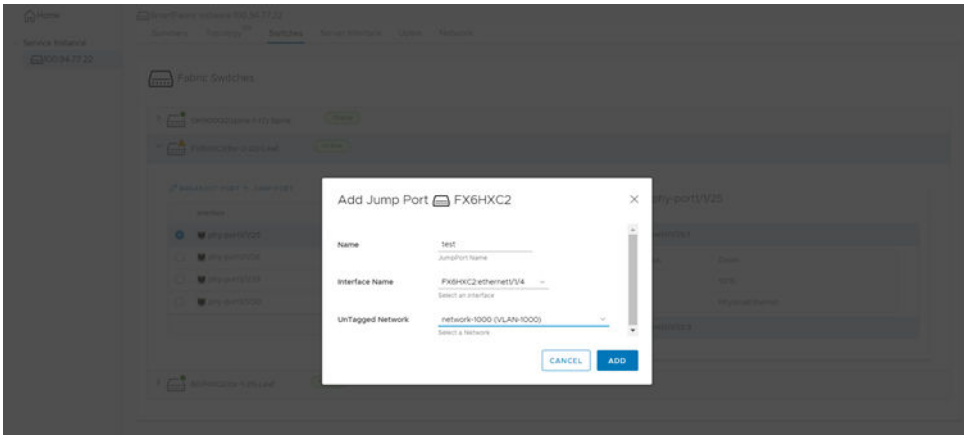

3. Jump ports add success; success/failure tooltip notification message displays.

Delete jump port

1. Select the switch to view the leaf switch properties to remove.

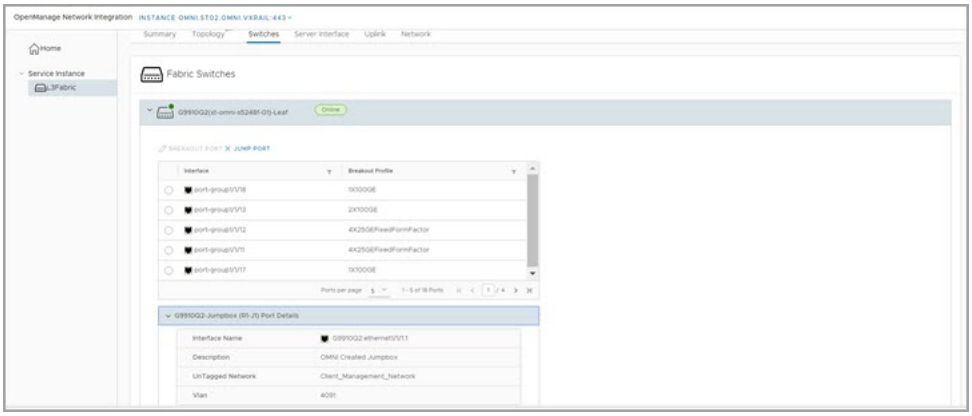

2. Click **Jump port**, then click **Delete**.

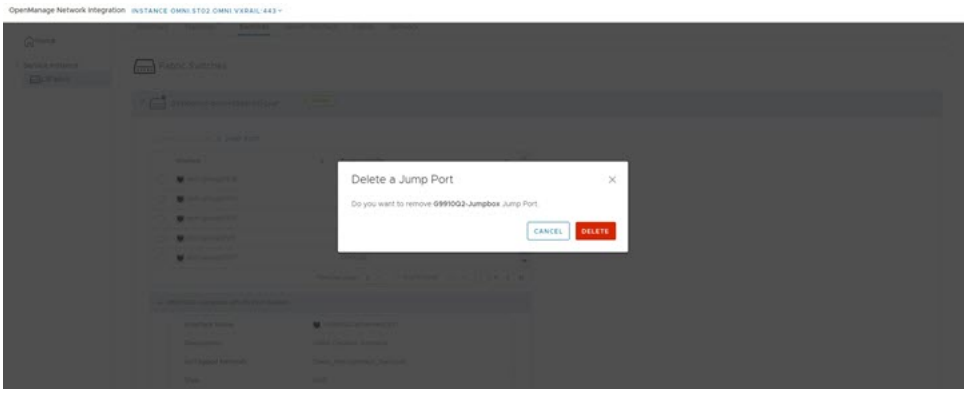

3. Jump port delete success; success/failure tooltip notification message displays.

Server interface

This information describes how to configure a SmartFabric server interface profile.

The Server Interfaces Profile List 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).

Create server interface profile

1. Select the **Service Instance**, then select **Server Interface**.

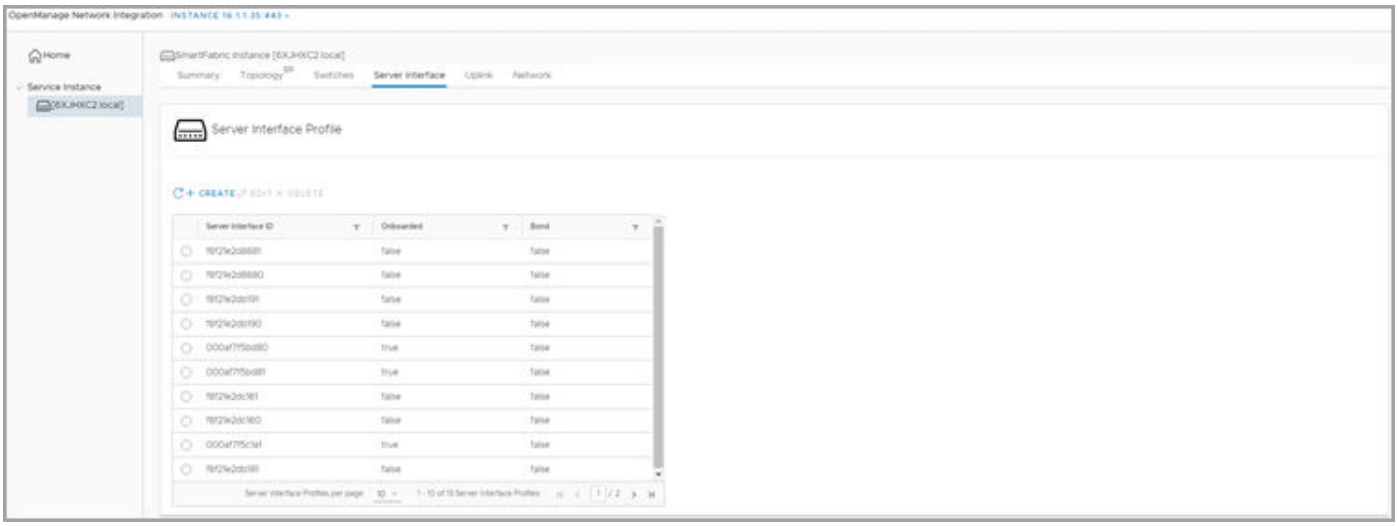

2. Click **Create** to create a Server Interface Profile and provide server interface ID, then select **New Server Profile**.

NOTE: System allows to 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 vmnic physical adapter MAC address without ":".

3. Select the **Server Profile Bonding Type**.

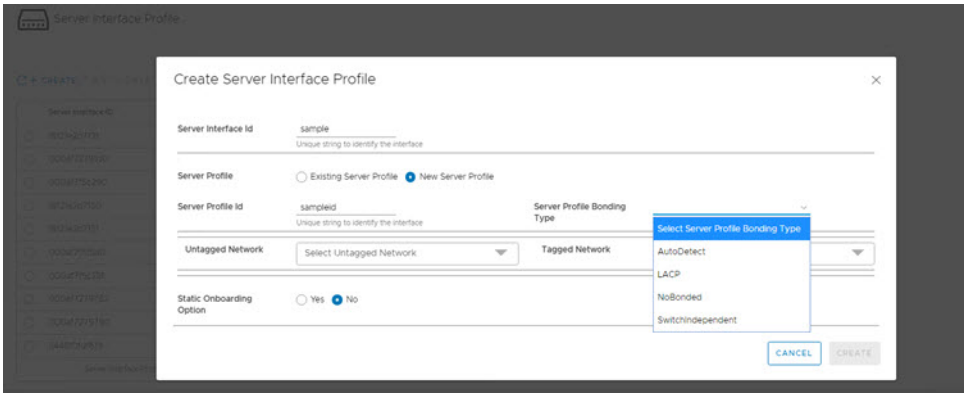

4. Select the **Untagged Network**, select the **Network**, then click **Create**.

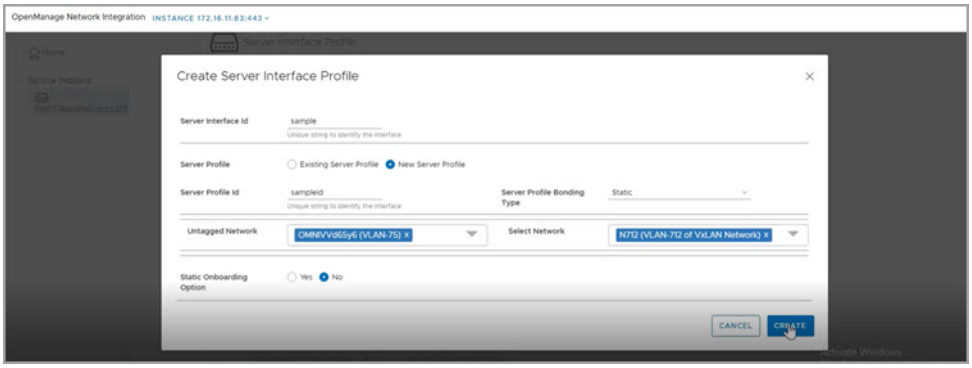

5. (Optional) Select **Yes** for the Static Onboarding Option, select Leaf Node and Interface (where the server interface is connected), select **None**, then click **Create**.

NOTE: You cannot modify the leaf switch and leaf switch interface in static onboarding server profile. If you want to modify, delete the existing server interface profile and re-add the new leaf switch and the leaf switch interface.

6. (Optional) Select **Yes** for the Static Onboarding Option, select Leaf Node and Interface (where the server interface is connected), select **eBGP**, enter the **ASN** and **IP address**, then click **Create**.

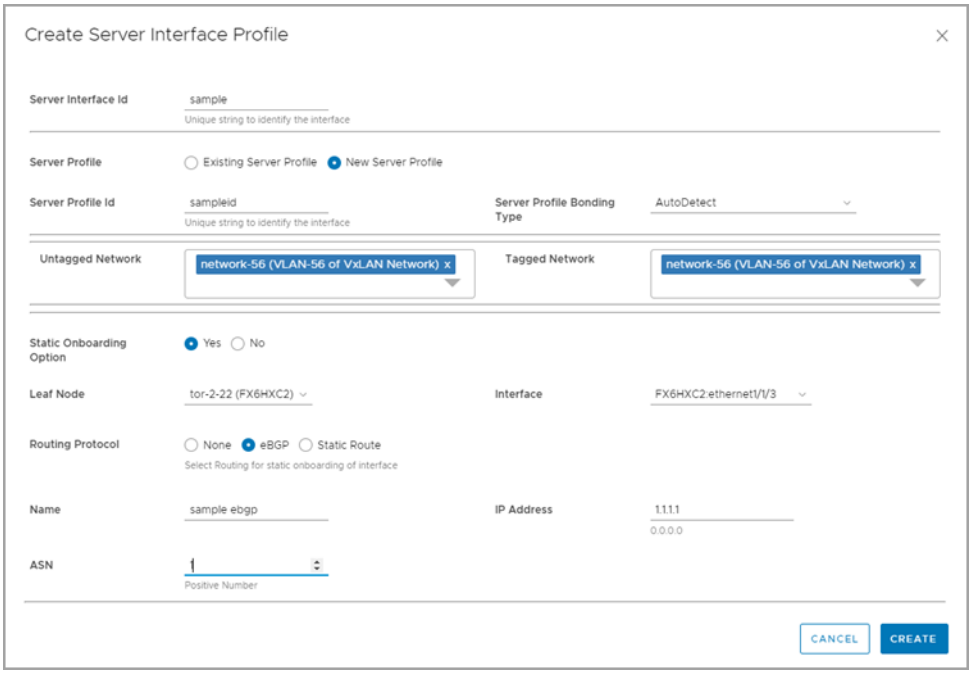

NOTE: Static onboarding option eBGP or Static route routing protocol used for NSX-T deployment.

7. (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**.

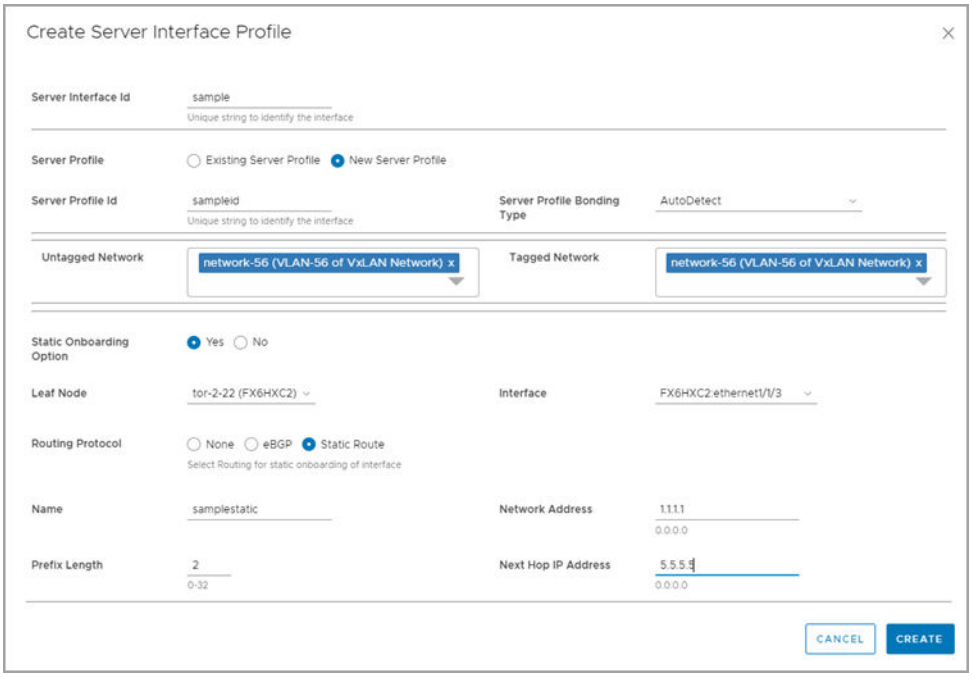

NOTE: You cannot delete any created server profile.

8. Server Profile and Service Interface creation successful; close the messages.

Edit profile

1. Select the Server Interface Profile, then click **Edit** to view the server interface information.

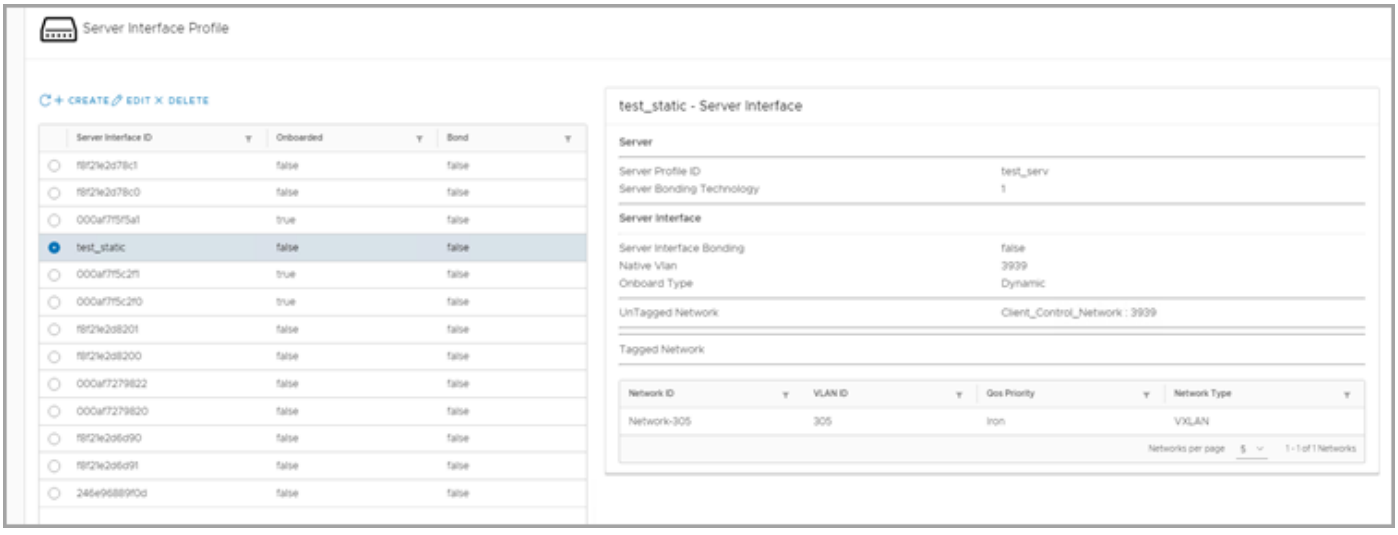

2. Select the **Untagged Network**, select the **Network**, then click **Edit**.

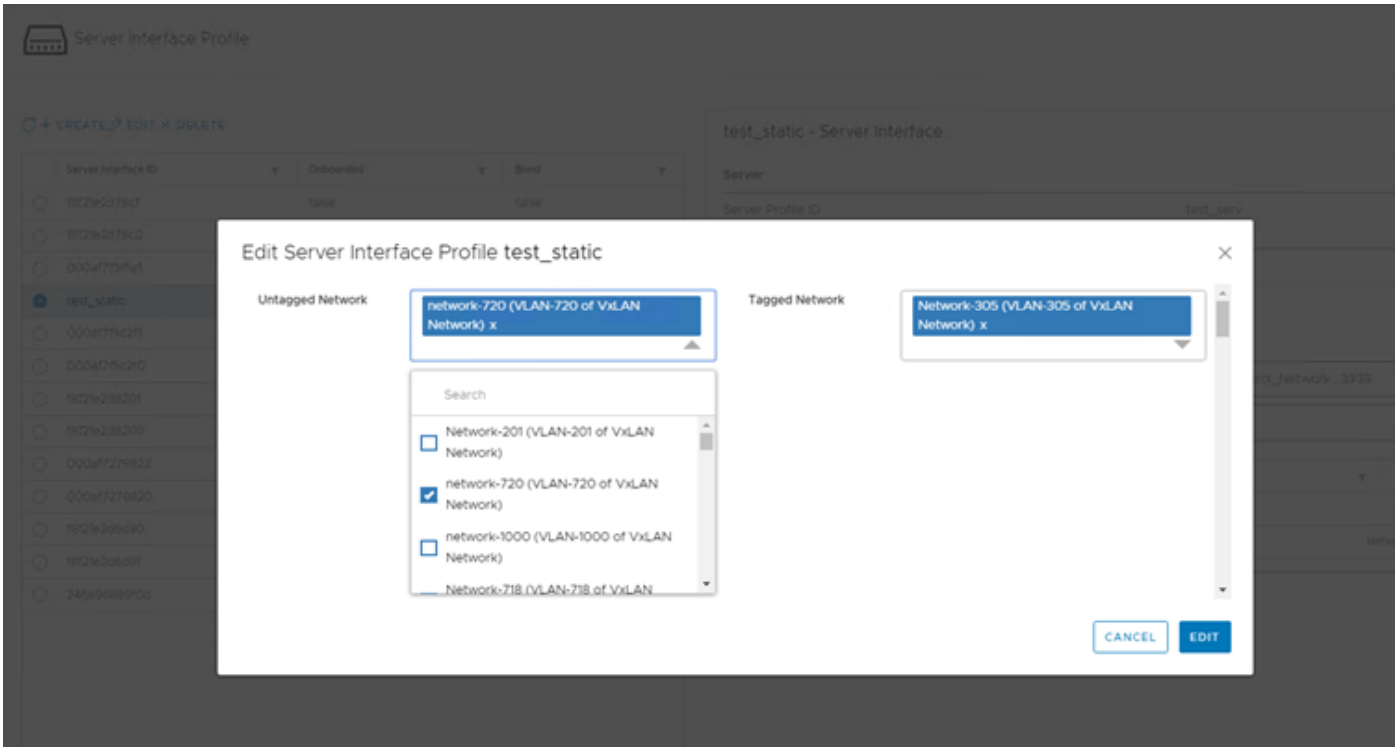

3. The Service Interface Profile successfully updated; close the message.

Delete profile

1. Select the Server Interface Profile from the display, then click **Delete**.

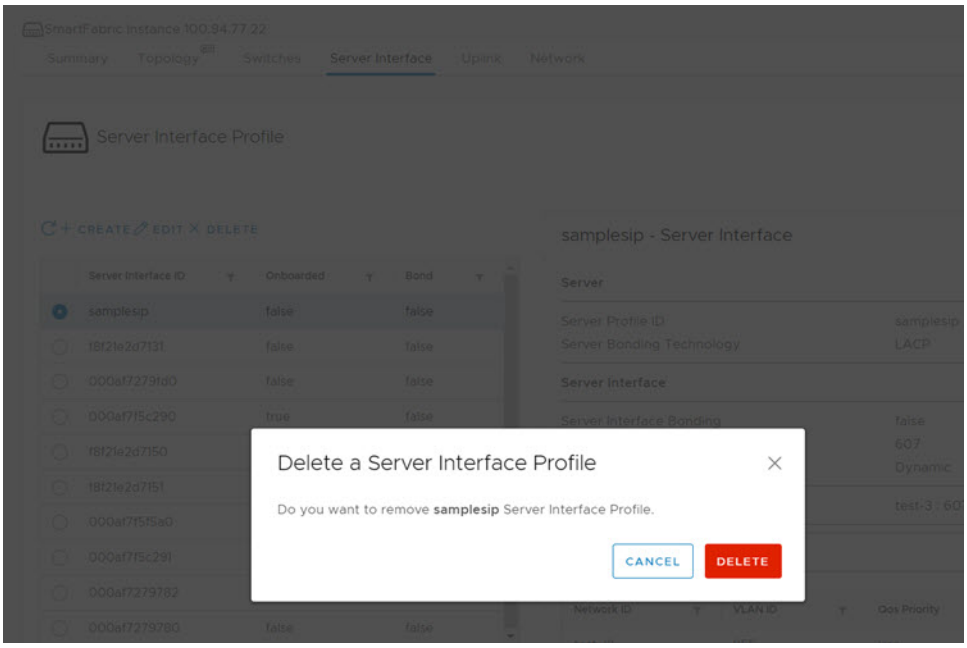

2. Click **Delete** to confirm removal of the Server Interface Profile.

Uplink

This information describes how to create, edit, or delete fabric uplinks. You can create uplinks on each fabric with available interfaces which are not part of an existing uplink, server connected ports, or part of a SmartFabric automation and 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 will be available for future use.

L2 uplinks

1. Select the **Service Instance**, then select **Uplink**.

Create user uplinks

You can create an uplink by selecting the fabric with a unique user-given name, then select the interfaces and networks through the wizard interface to create a user uplink.

1. Select **Uplink**, then click **Create**.

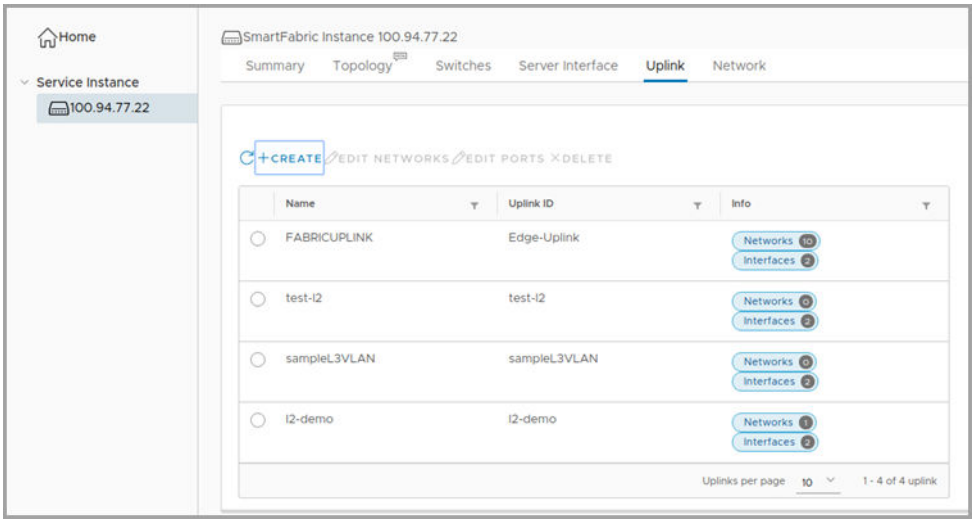

2. Enter the Uplink Name, an optional description, then click **Next**.

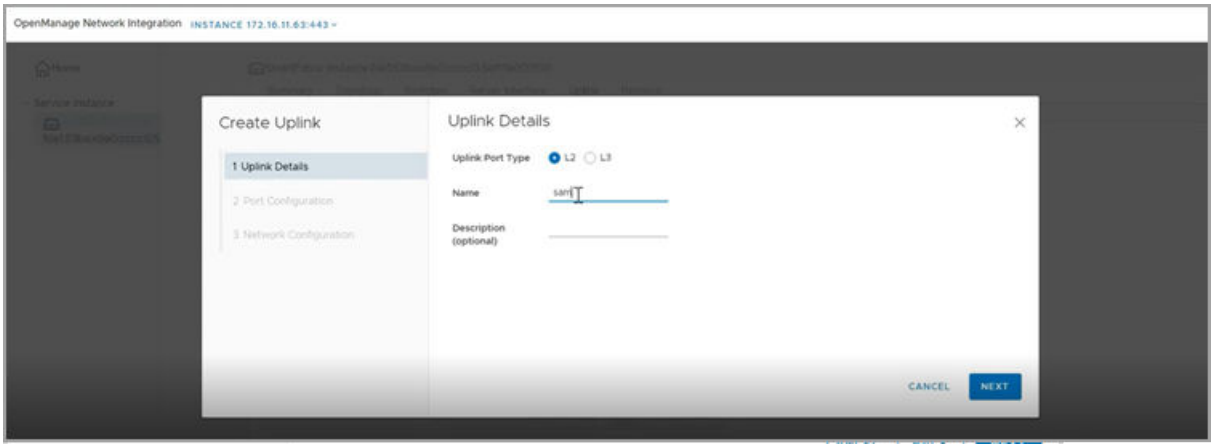

3. Select the rack to create the uplink on, select the interfaces, the LAG mode (LACP or Static), then click **Next**.

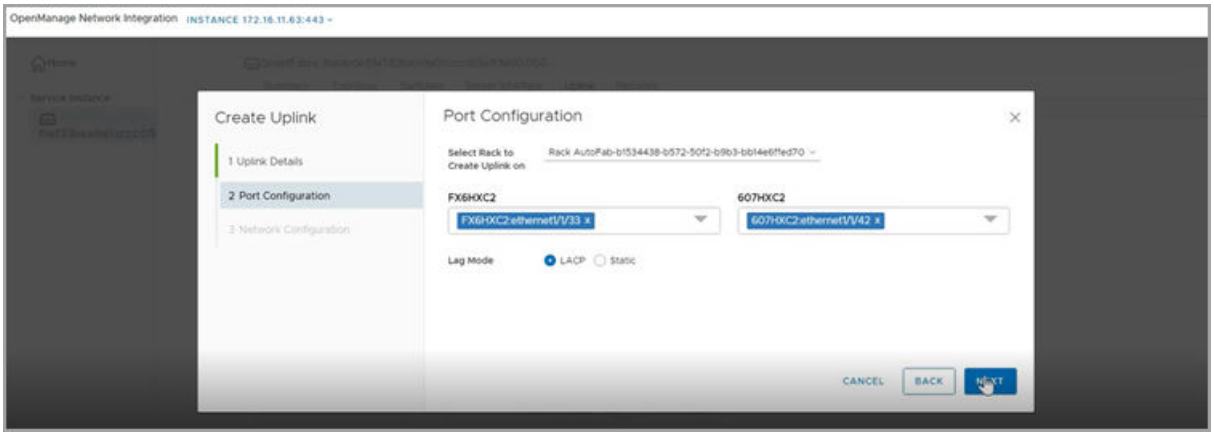

4. Select the untagged network, select the OMNI network, then click **Finish**.

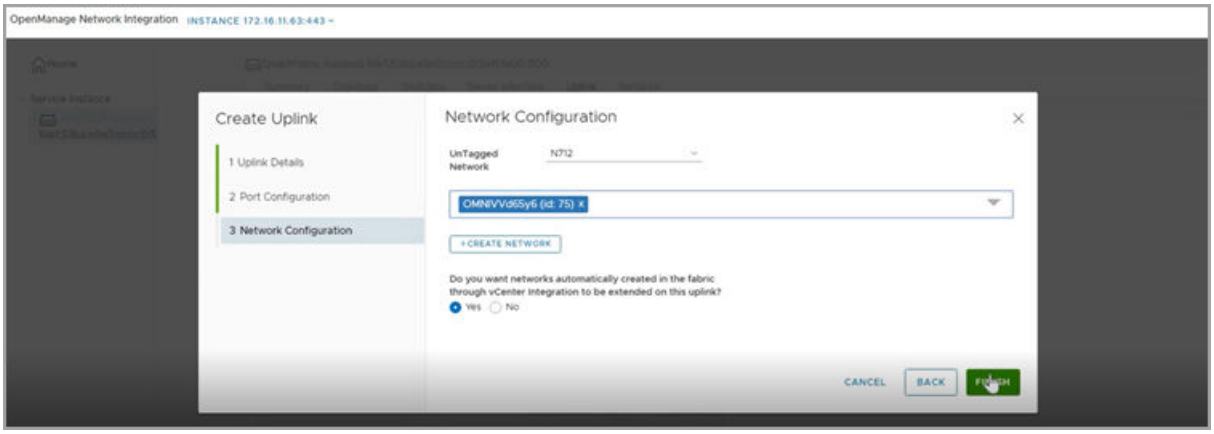

5. User uplink creation successful; close the message.

Edit user uplink

1. Select the **User Uplink**, then click **Edit Networks**.

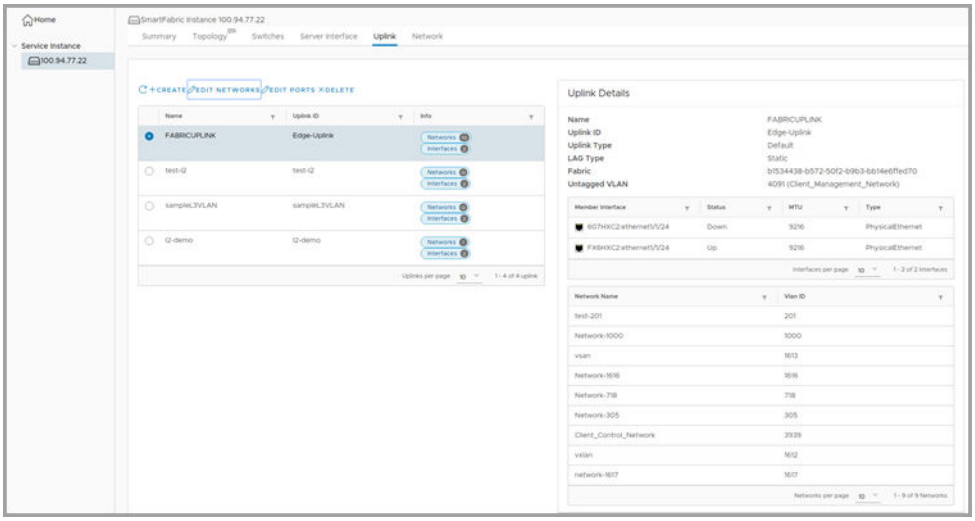

2. Select the Untagged Network to edit, select the OMNI network, then click **Update**.

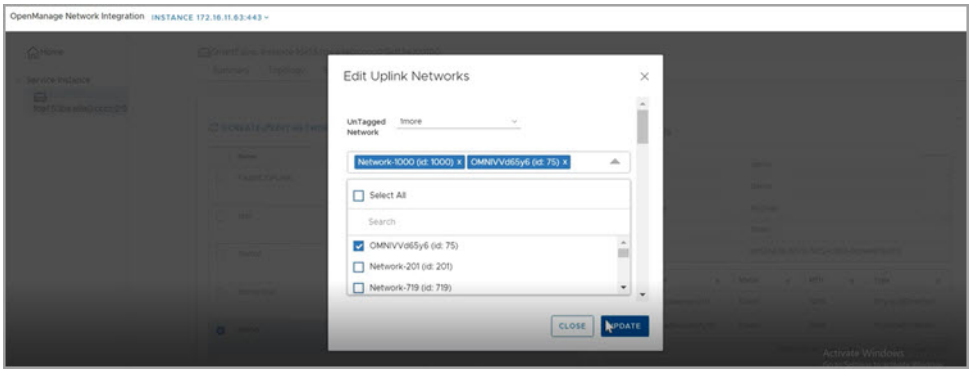

3. Click **Edit Ports**.

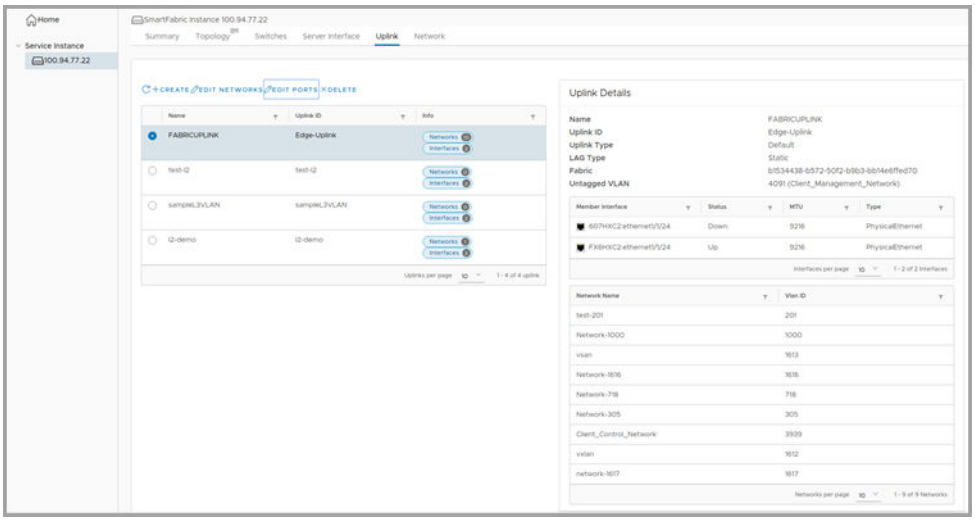

4. Verify the uplink interfaces, make changes as necessary, then click **Update**.

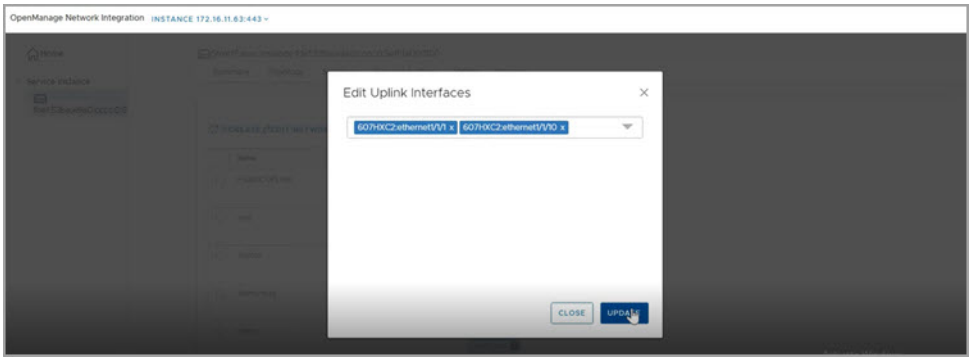

5. Editing of uplink interfaces successful; close the message.

Delete user uplink

NOTE: You can only delete user-created uplinks. Default system uplinks cannot be deleted.

1. Select the uplink to delete, then click **Delete**.

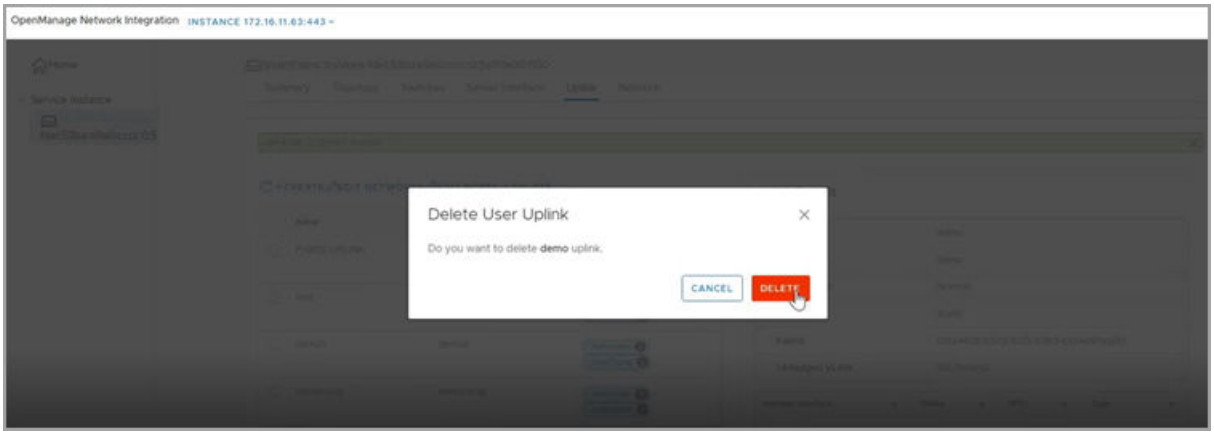

2. Click **Delete**.

L3 uplinks

L3 VLAN uplink

1. 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**.

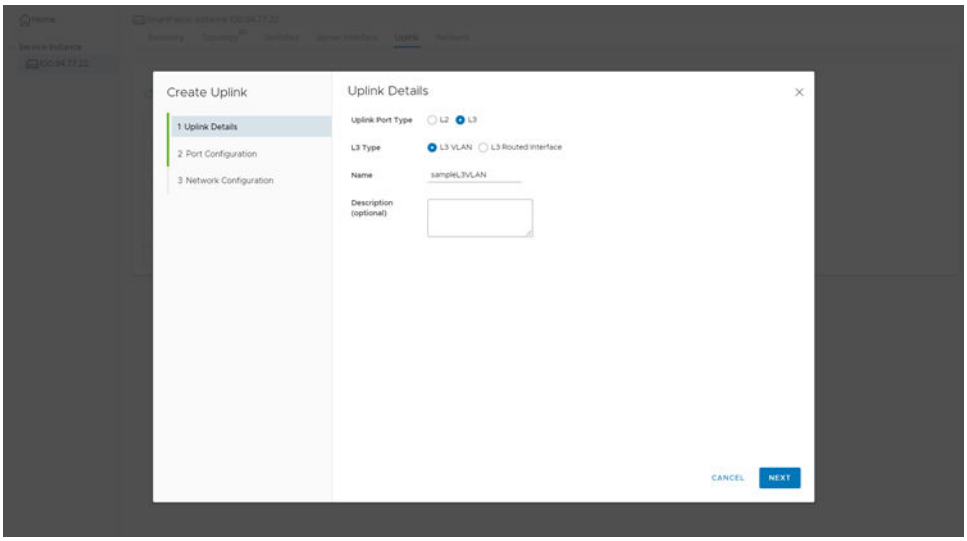

3. Select the **Switch group** (Leaf or Spine), select the **rack** to create the uplink on, select the **interfaces**, select **LACP** for the LAG mode, then click **Next**.

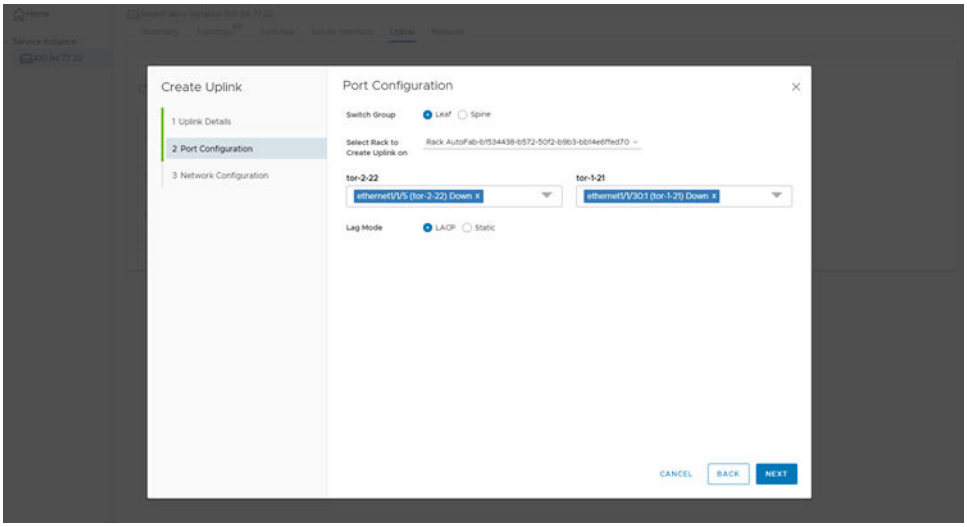

4. Select **UnTagged** network, select the **OMNI network**, enter an optional description, select either eBGP or Static Route for the routing protocol, enter the routing policy information, then click **Finish**.

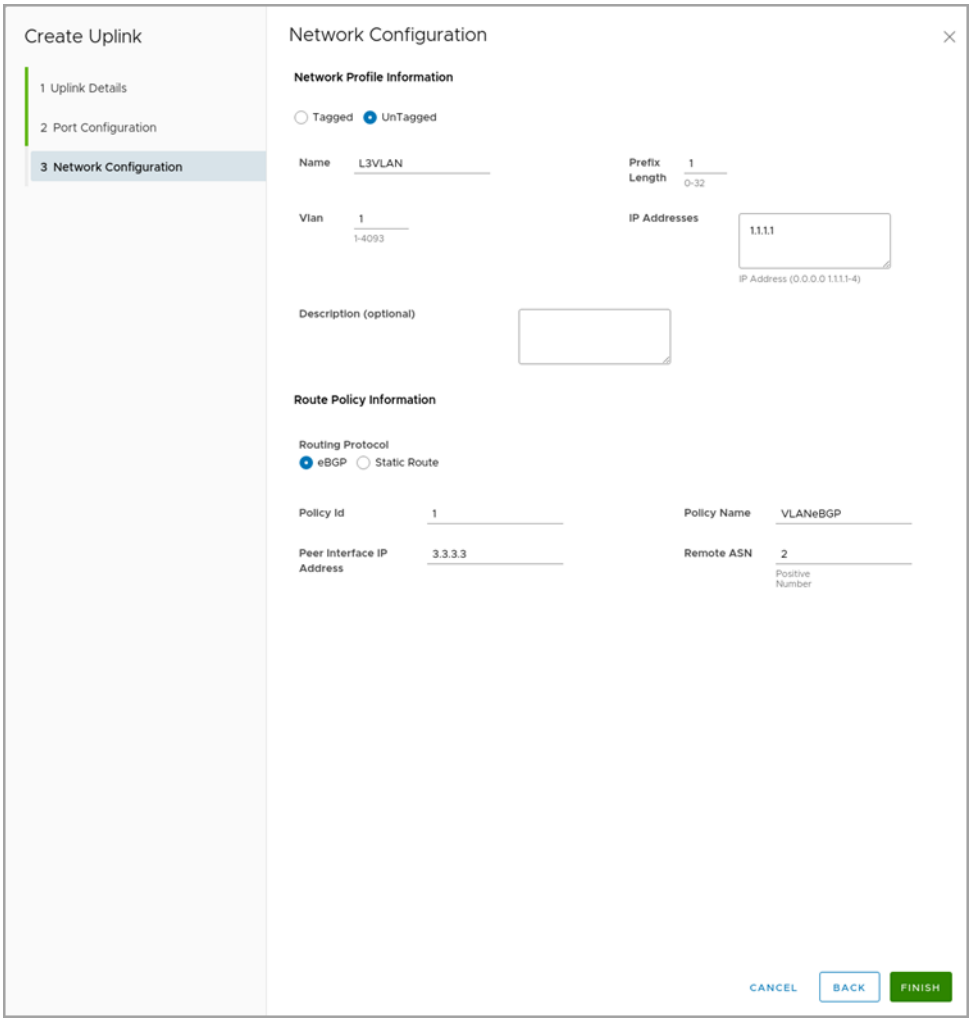

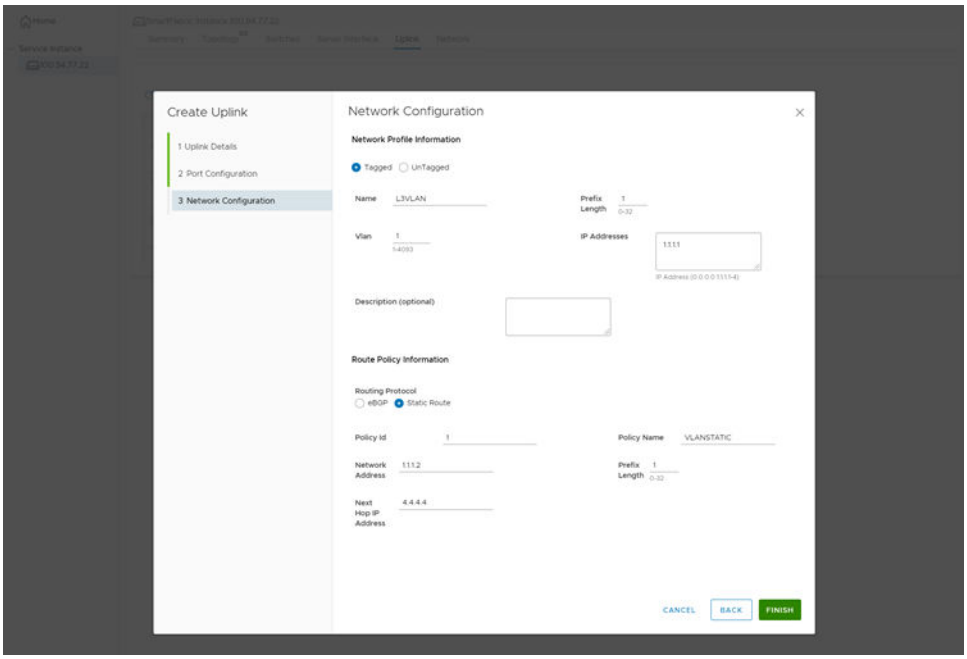

5. A route is created associated with the nodes that are configured in the port configuration. Uplink creation success; close the message.

L3 routed interface uplink

1. Click **Create**.

2. Select **L3 routed interface**, enter the **Uplink name**, and optional description, then click **Next**.

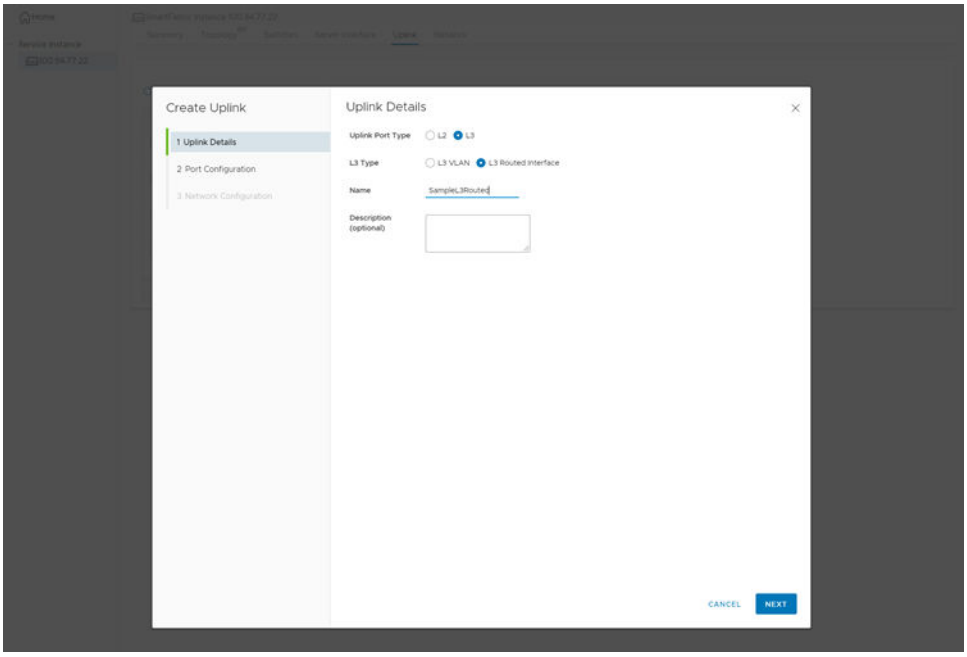

3. Select the Switch group (Leaf or Spine), the **rack** to create the uplink on, select the **interfaces**, then click **Next**.

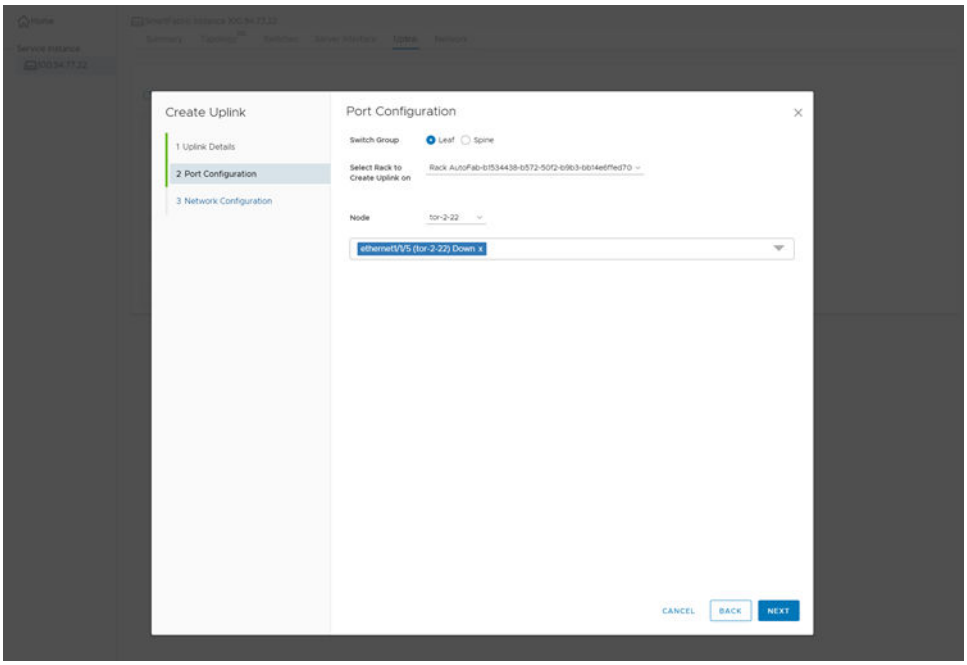

4. Enter the network profile information and routing policy information for the uplinks, then click **Finish**.

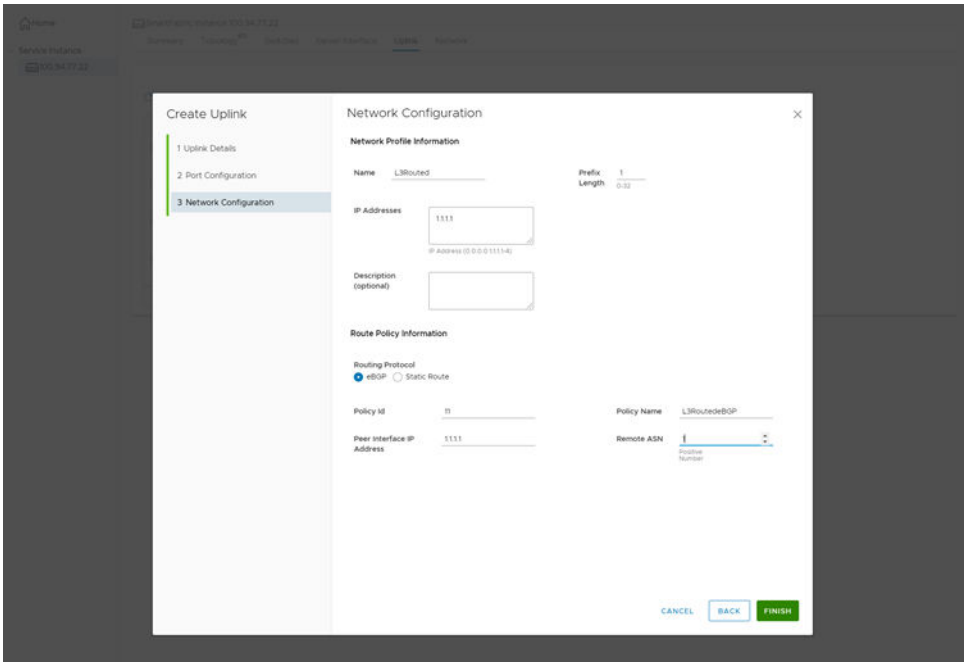

5. L3 routed uplink creation success; close the message.

Network

This information describes how to configure networks, and routing configuration. You can configure three types of networks including VxLAN networks (for L2 and L3 profiles), VLAN networks (for L2 and L3 profiles), and L3 routed interfaces (for L3 profiles only).

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. You should add *Manual* networks to uplinks using the OMNI portal if needed. The OMNI VLAN automation process uses *Manual* networks for `ServerInterfaces`. We recommend not using the OMNI portal to create a network, if you are using those VLANs for the OMNI registered vCenter `PortGroup`. OMNI automation manages those VLANs/networks by itself. For complete information, see OMNI vCenter integration.

VxLAN networks

Create VxLAN network

1. Select **Networks > VxLAN Networks**, then click **Create**.

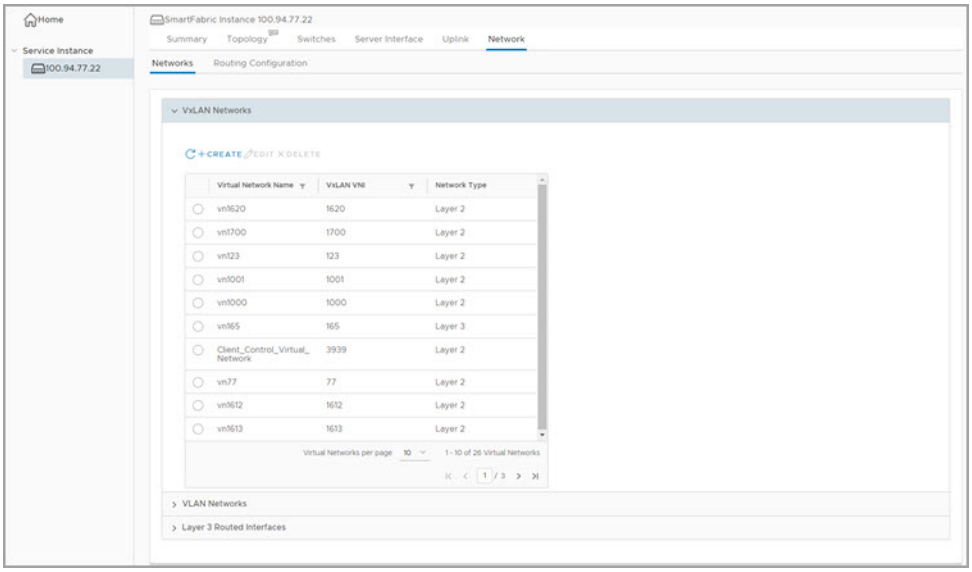

2. Verify **Layer 2** is selected as the Virtual Network Type, enter the **Virtual Network Name**, enter a value between 1 and 16777215 for the VxLAN VNI, enter a value between 1 and 4093 for the VLT VLAN ID, enter an optional description, then click **Create**.

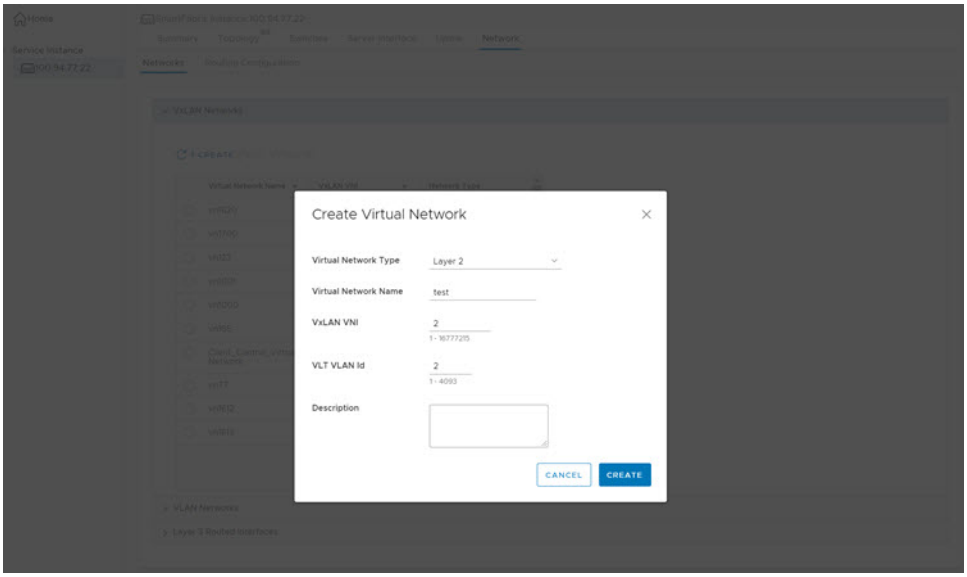

3. Virtual network creation successful; close the message.

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.

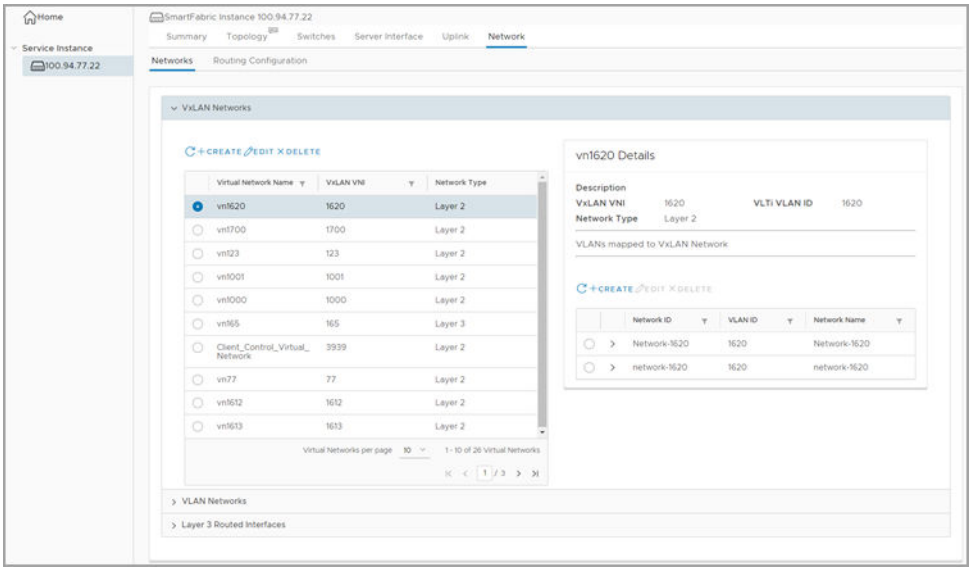

Edit VxLAN network

1. Select the Virtual Network Name, then click **Edit**.

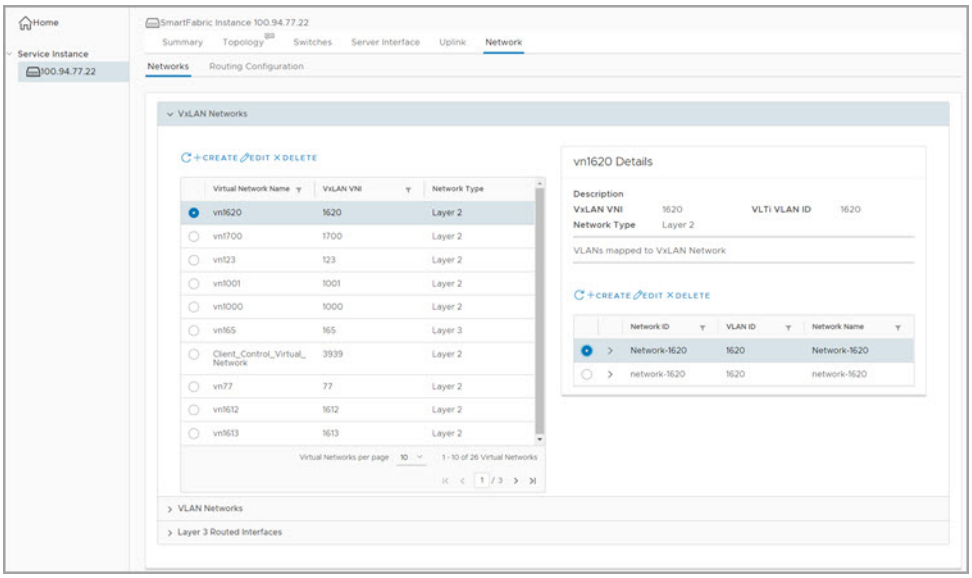

2. Modify the Virtual Network Type, then click **Edit**.

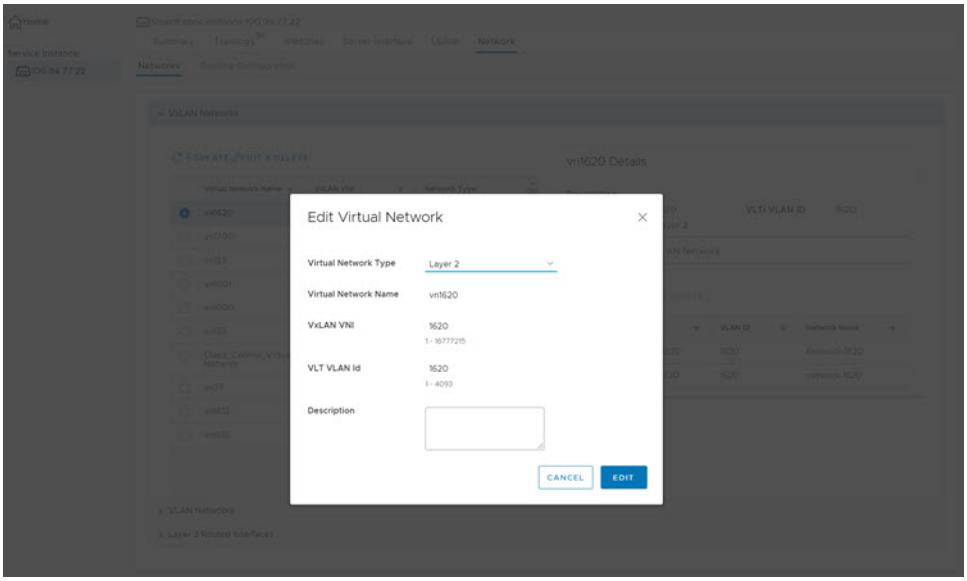

3. Modify the Prefix, Gateway IP Address, IP address, then click **Edit**.

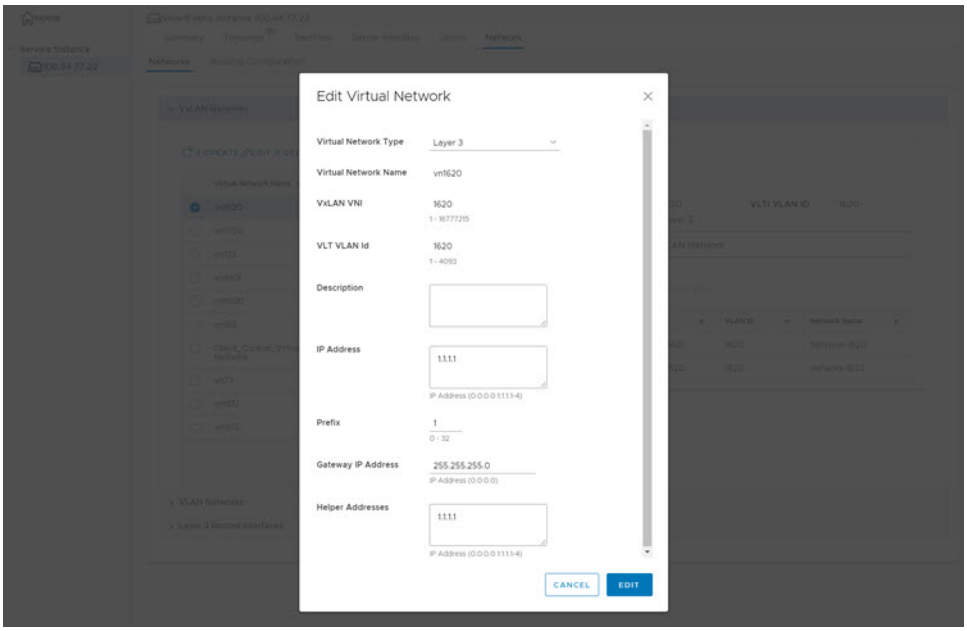

4. Virtual network edits successful; close the message.

Delete VxLAN network

1. Select the **Virtual Network Name**, select the **Network**, then click **Edit** on the right.

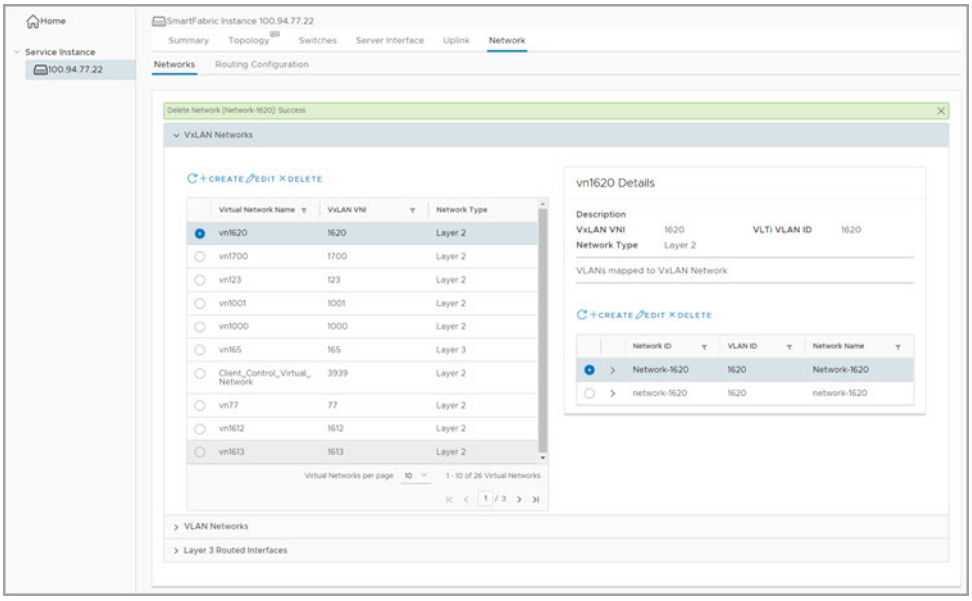

2. Click **Delete** to remove the network.

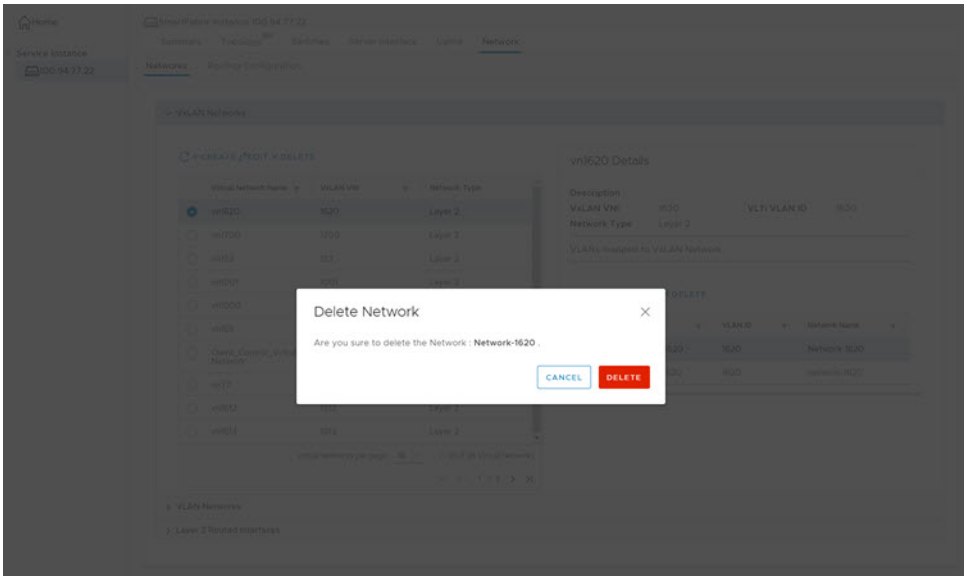

3. Delete network success; close the message.

VLAN networks

Create network

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

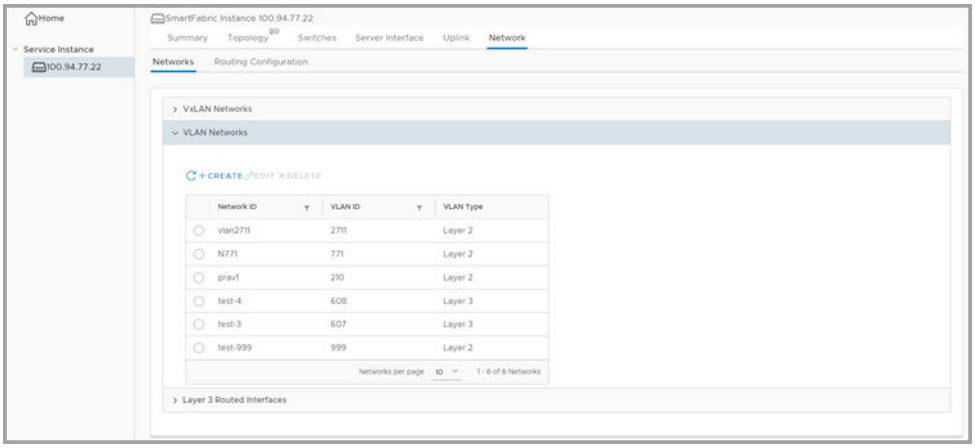

2. Verify **Layer 2 VLAN Network** is selected as the Network Type, enter the **Network ID**, **Network Name**, enter 1 to 4093 for the VLAN, enter an optional description, then click **Create**.

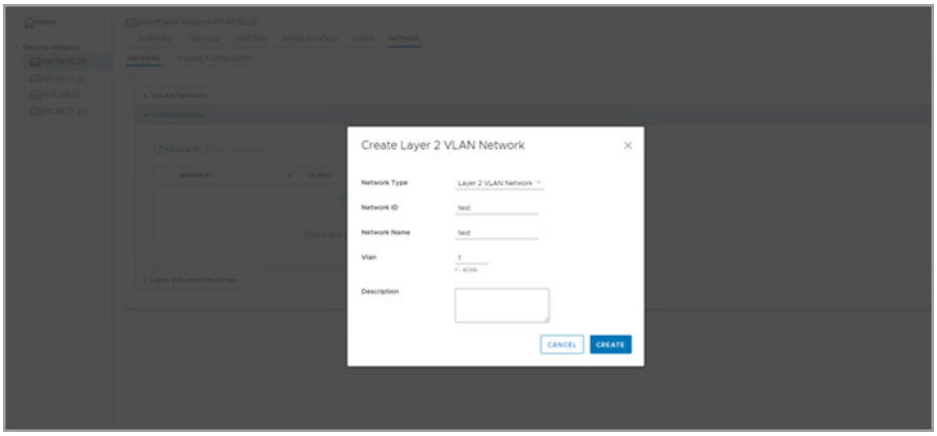

3. VLAN network created; close the message.

Edit network

1. Select the Virtual Network Name, then click **Edit**.

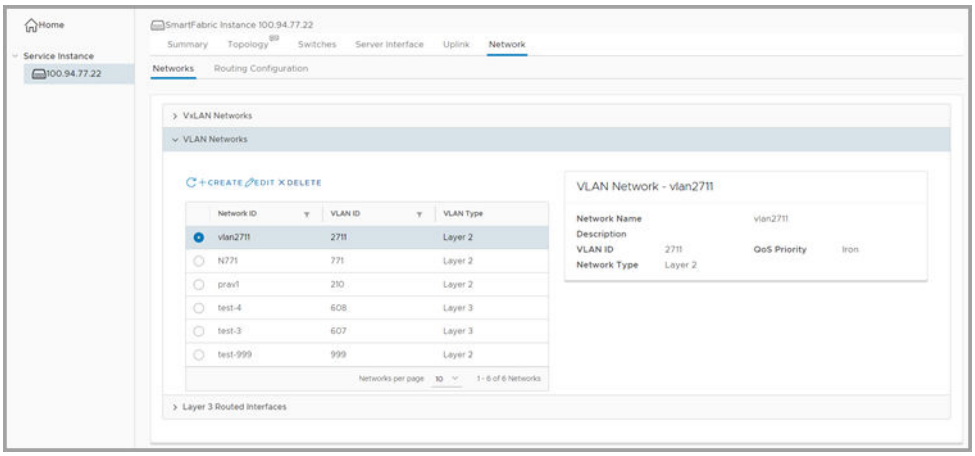

2. Verify the Network Name, make changes as necessary, then click **Edit**.

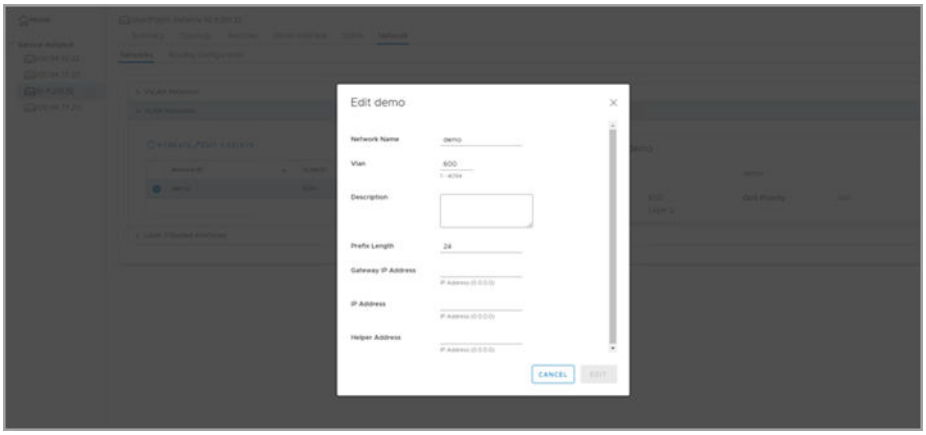

3. Edit network success; close the message.

Delete network

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

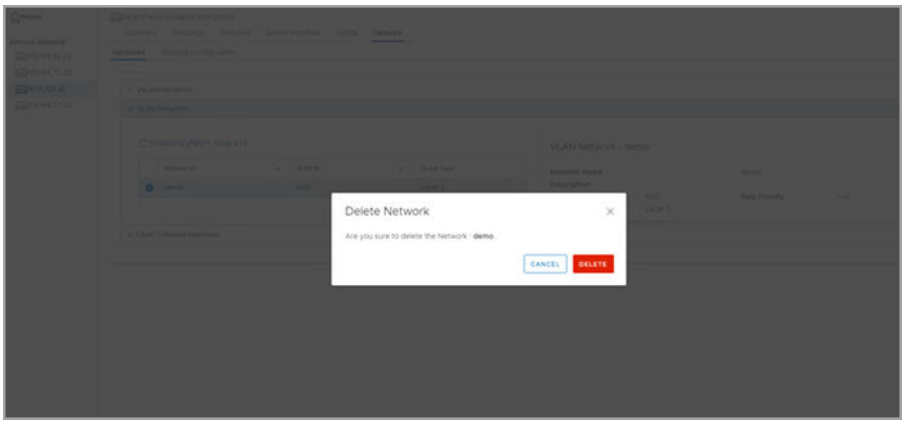

2. Click **Delete** again.
3. Network delete success; close the message.

L3 routed interfaces

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

Create network

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

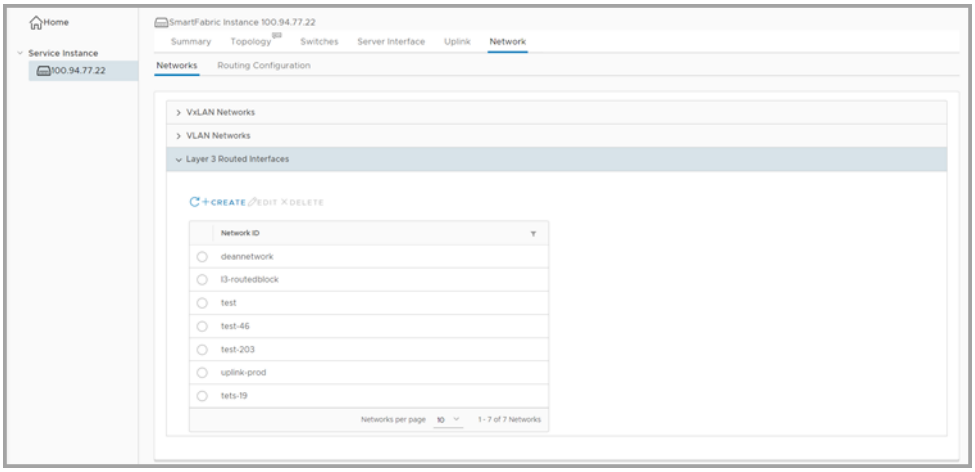

2. Enter the **Network ID**, **Network Name**, select the **Prefix Length**, select the **IP Address**, enter an optional description, then click **Create**.

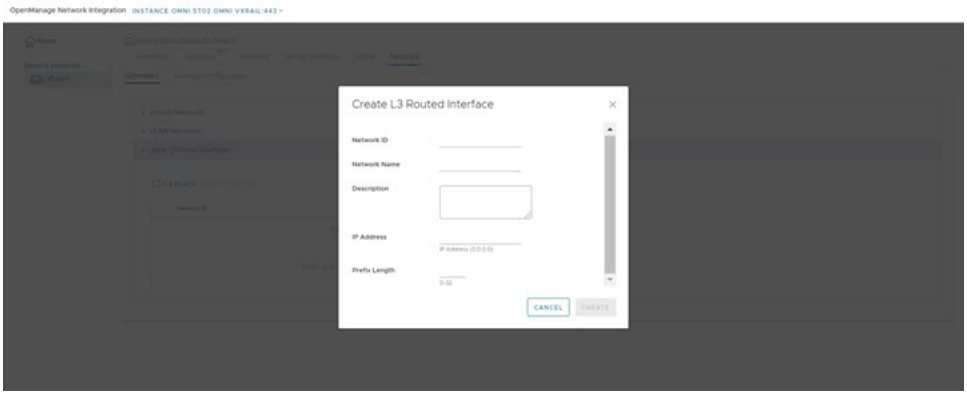

3. Network create successful; close the message.

Edit network

1. Select the Network ID, then click **Edit**.

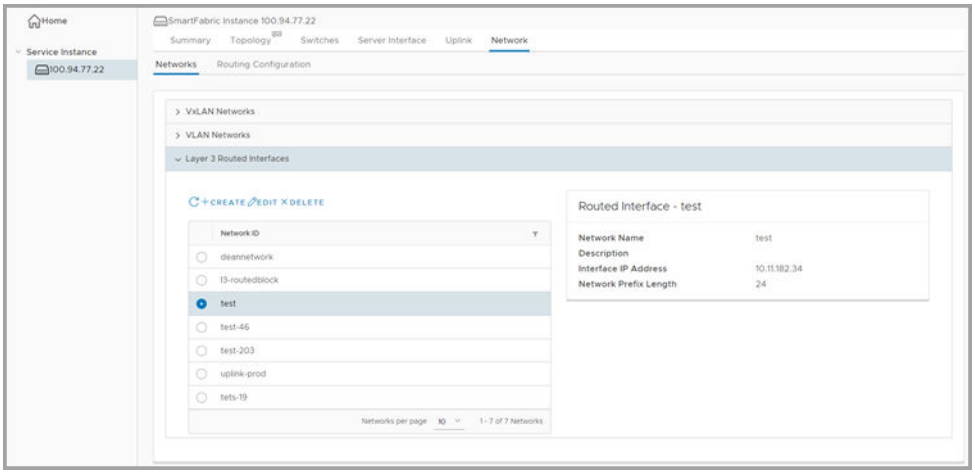

2. Make changes as necessary, then click **Edit**.

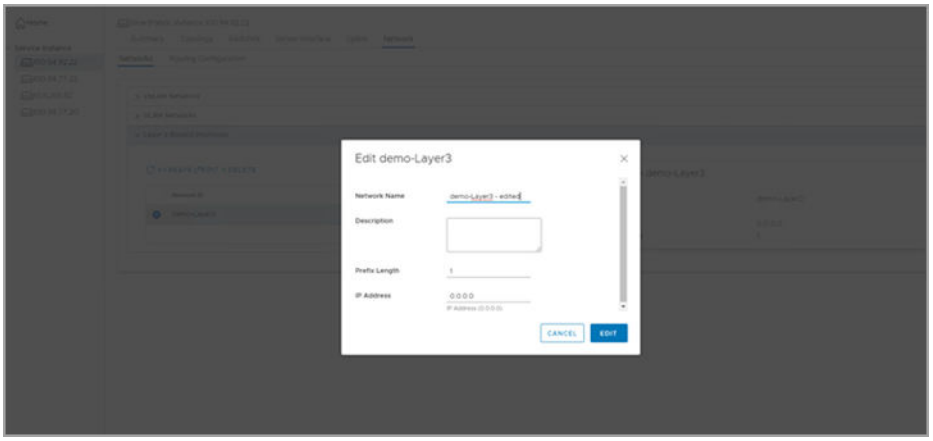

3. Network edit success; close the message.

Delete network

1. Select the network to remove, then click **Delete**.

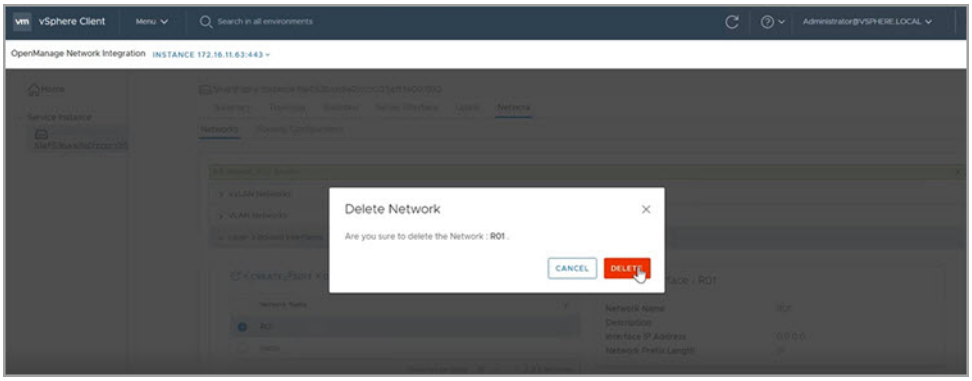

2. Delete network success; close the message.

Routing configuration

This information describes static routes, and eBGP peer configuration.

NOTE: Creation of routes is only allowed in Uplinks or the Server Interface Profile.

Delete static route

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

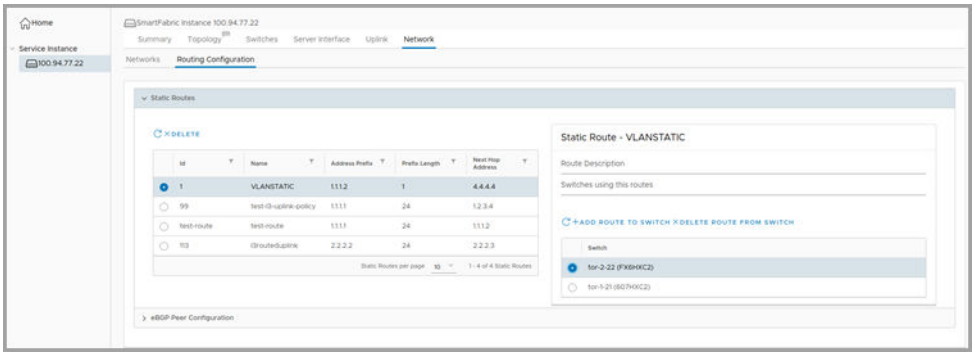

2. Click **Delete Route**.

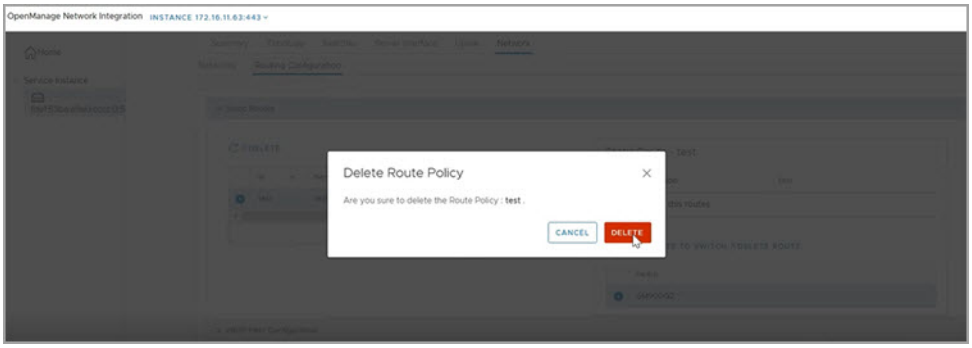

3. Delete static route success; close the message.

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.

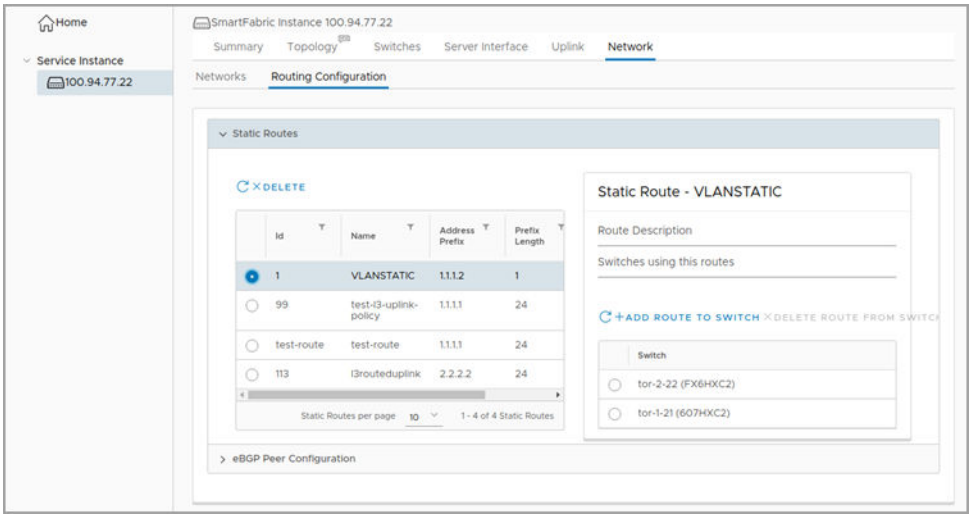

Add route to switch

1. Select **Routing Configuration > Static Routes**, select a static route, then click **Add Route to Switch**.

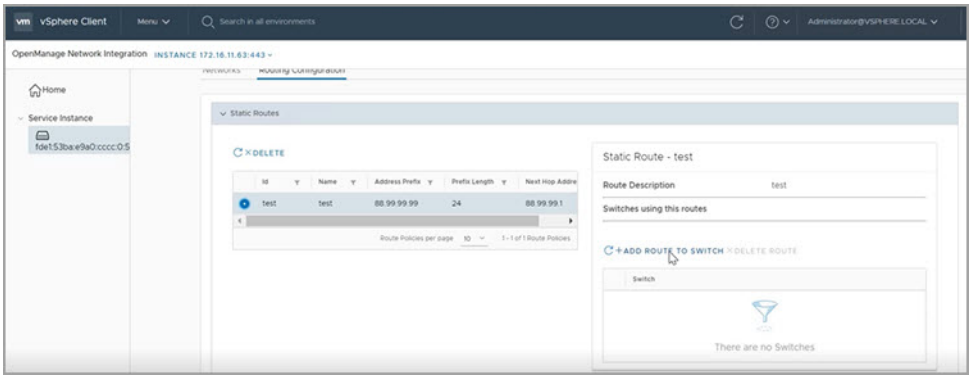

2. Select the switch to map to this route, then click **Add**.

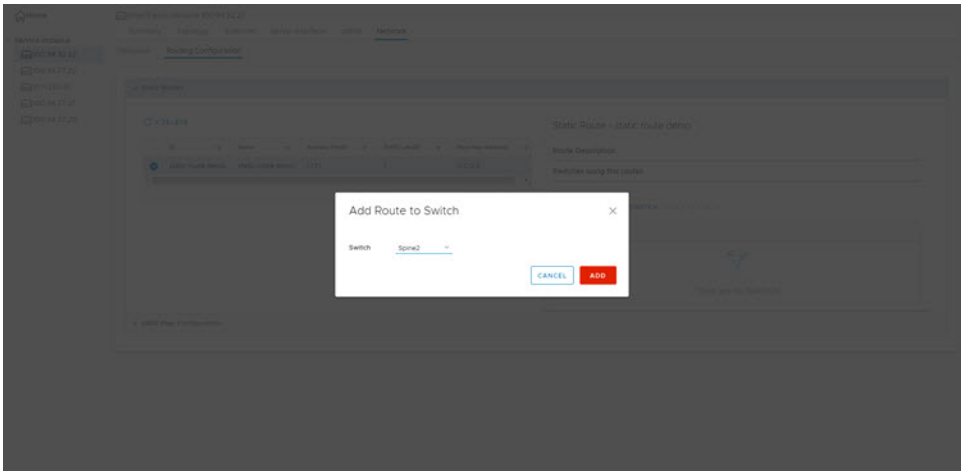

3. Route added to switch success; close the message.

Delete route from switch

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

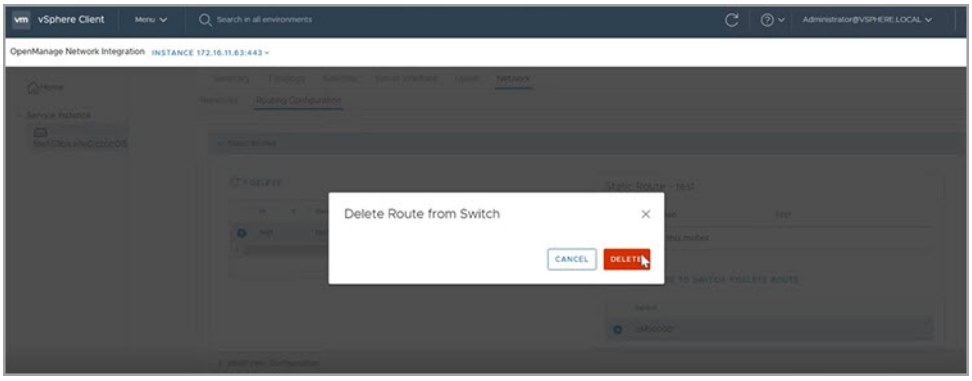

2. Click **Delete** to remove the route from the switch.
3. Delete route policy successful; close the message.

Delete eBGP route

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

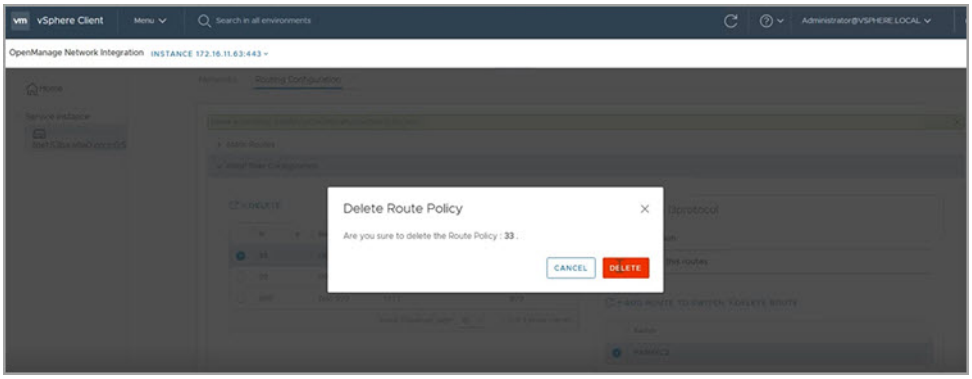

2. Delete route policy success; close the message.

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.

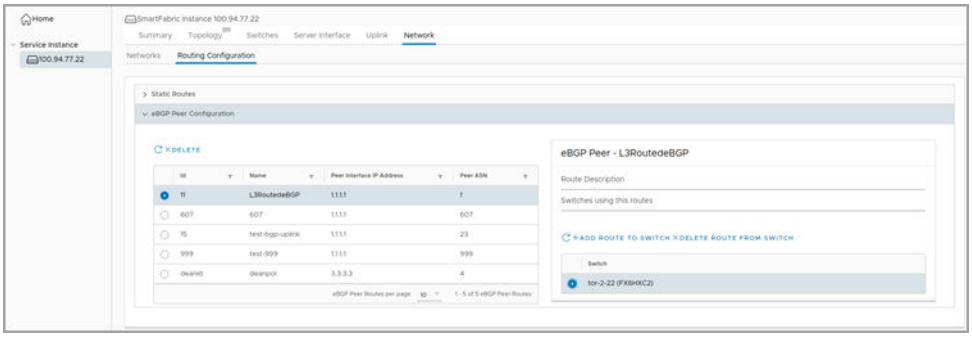

Add eBGP route to switch

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

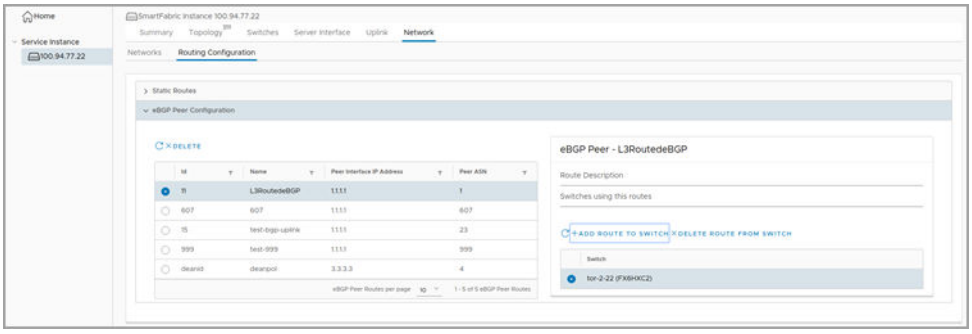

2. Select the switch, then click **Add**.

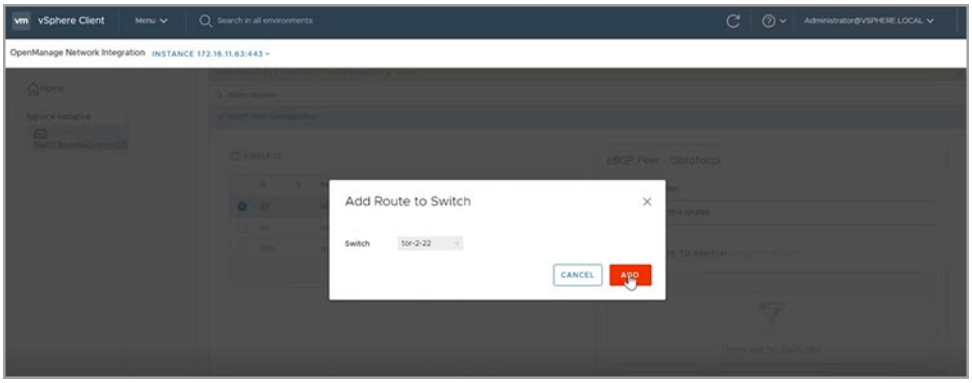

3. Route switch add success; close the message.

Delete eBGP route from switch

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

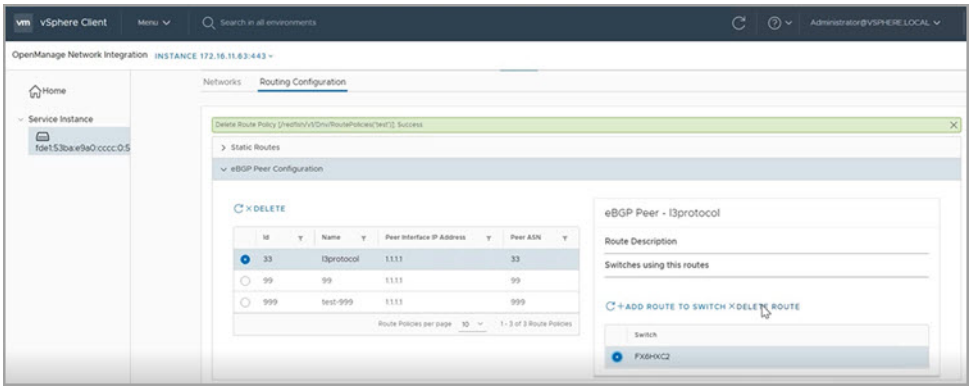

2. Click **Delete** to remove the route from the switch.

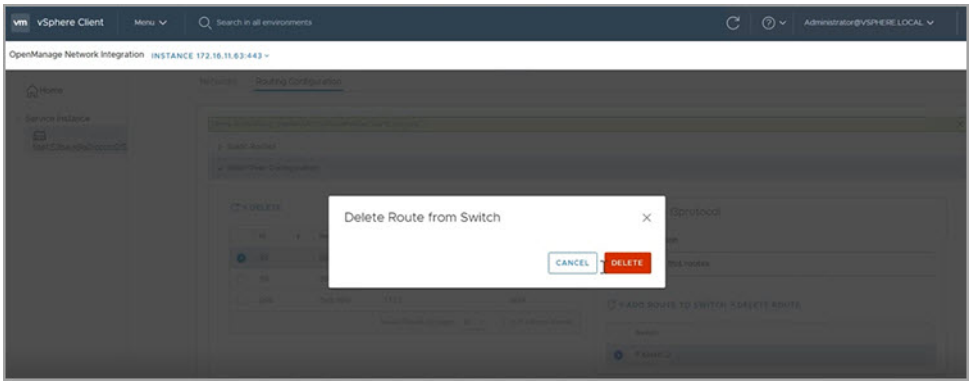

3. Delete route success; close the message.

Host network inventory

This information describes how to 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.

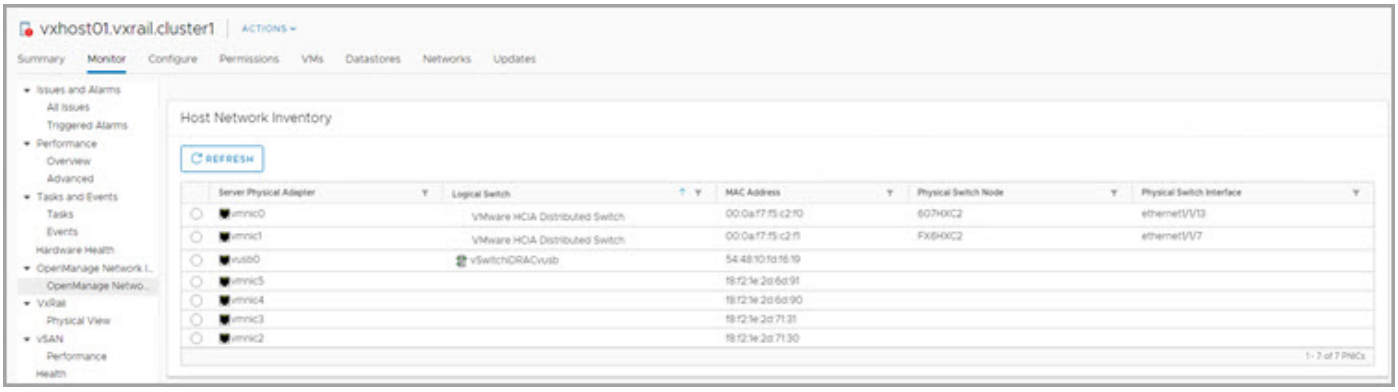

- Physical adapter — physical network adapter name
- Virtual switch — name of switch the physical adapter is connected to
- MAC address — MAC address of the physical adapter
- Physical switch — physical switch that is connected to the fabric
- Physical switch interface — physical switch port this server network adapter is wired to

Logical switch

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

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

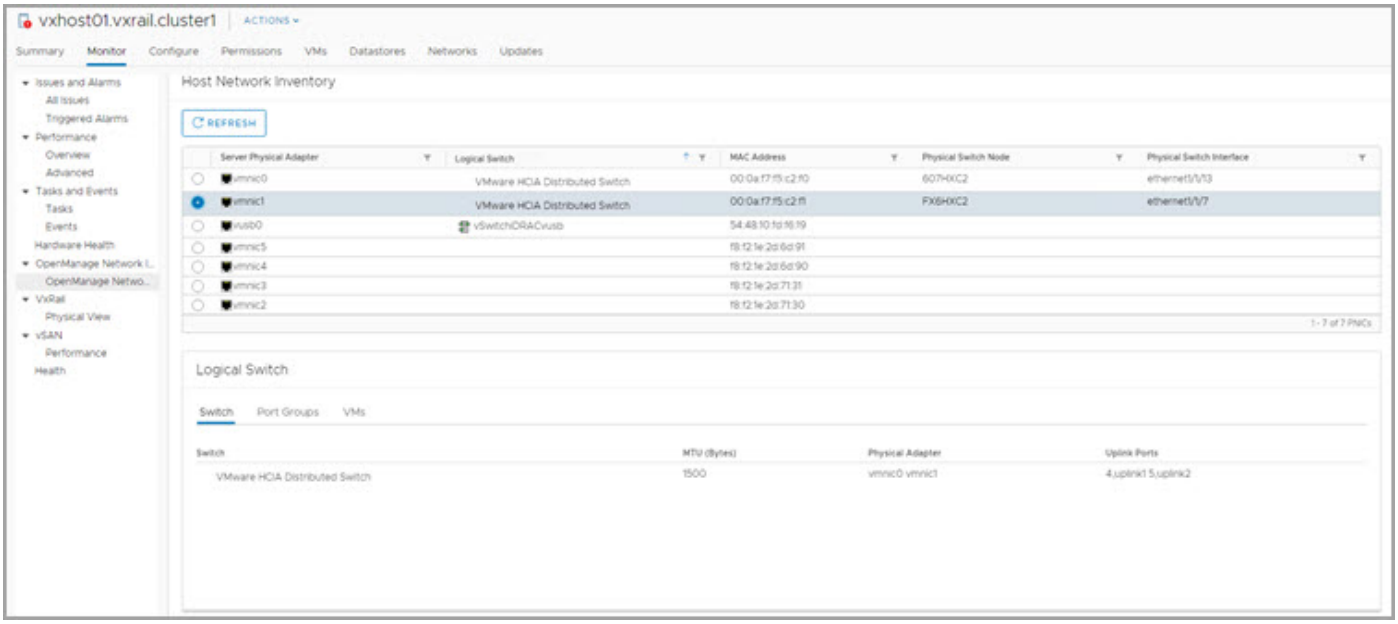

- Port groups tab — includes the name of port groups, and VLAN IDs for each port group

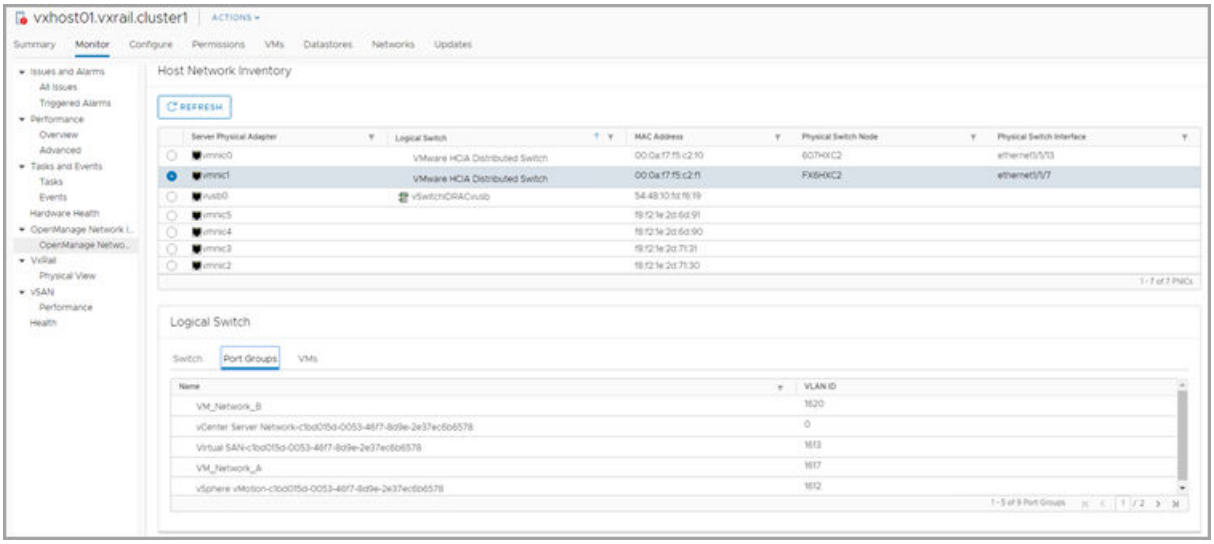

- VMs tab — includes the name of VMs of that host that is connected to a single virtual switch

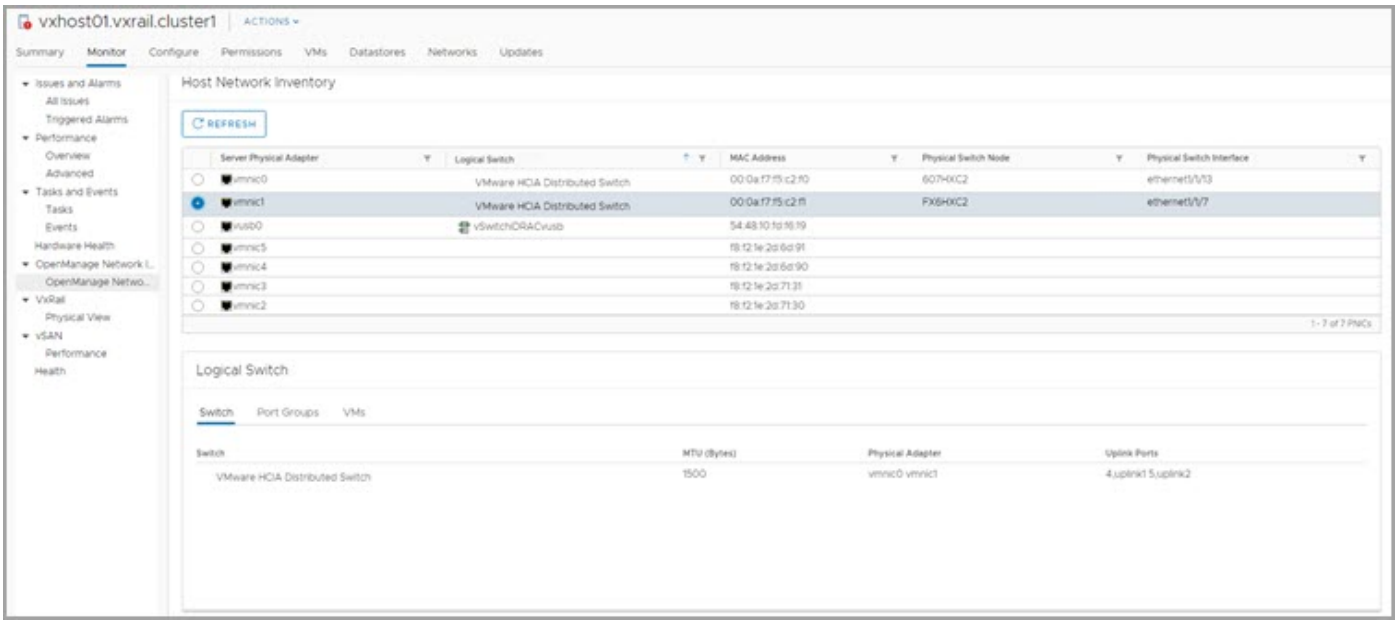

Physical switch

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

- Onboard interface

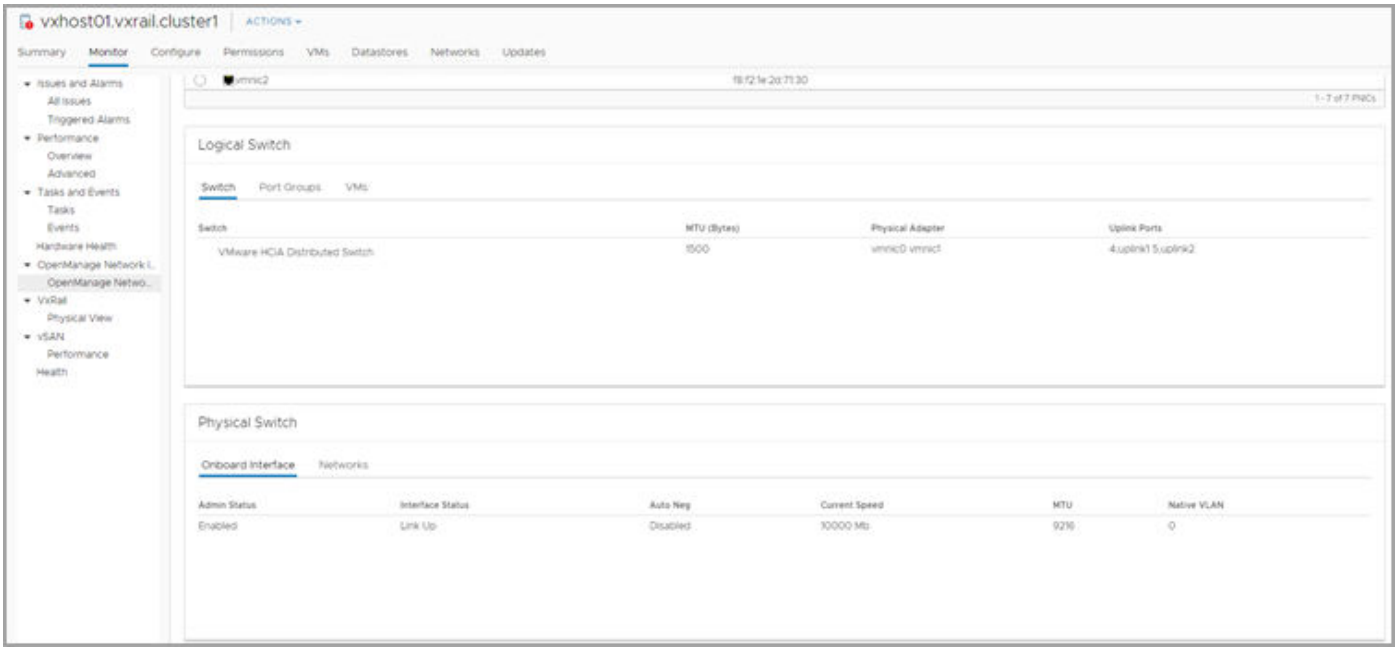

- Admin Status — configured state of the physical interface
- Interface Status — current operations state of the physical switch port
- Auto Neg — negotiation status of the physical interface
- Current Speed — current operational speed of the physical interface
- MTU — maximum transmitting unit configured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

- Networks

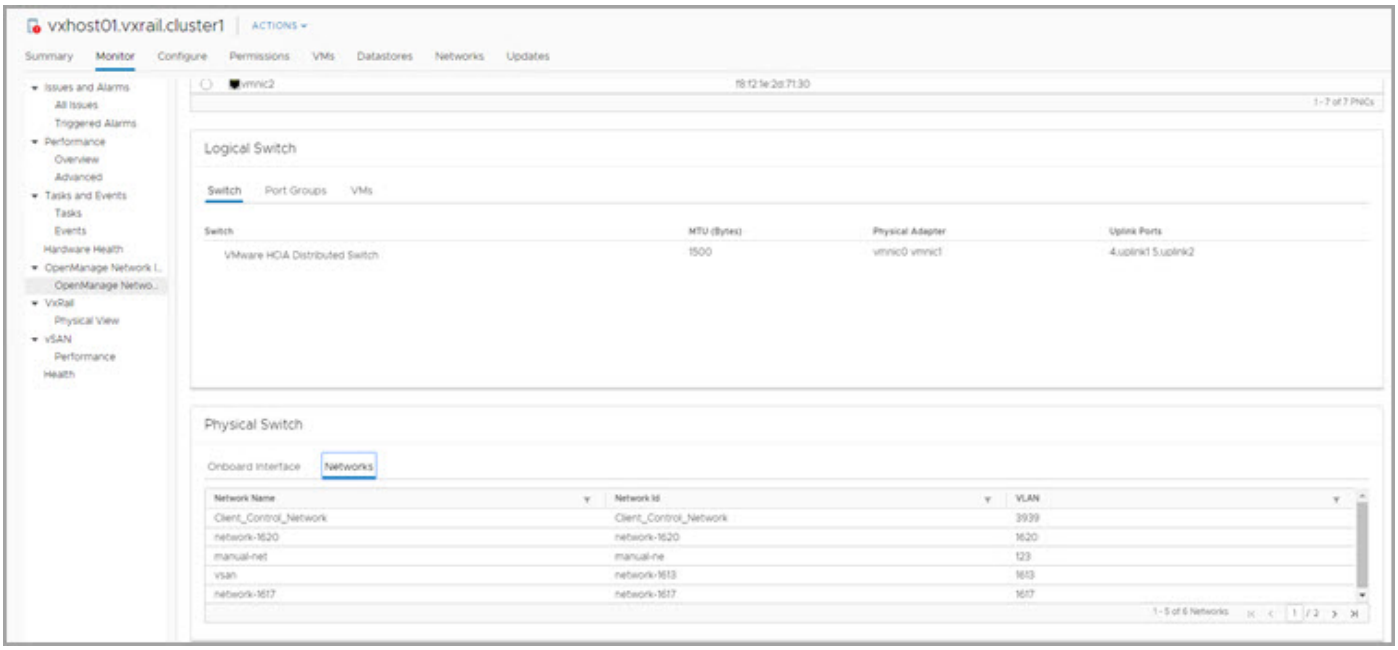

- Network Name — name of the VLAN network
- Network ID — unique identifier of the fabric network
- VLAN — tagged VLAN of the switch port

8

Network fabric management

This information explains common operations including upgrading of the SmartFabric solution including VxRail nodes, SmartFabric OS10, SmartFabric vCenter appliance, and switch replacement.

Topics:

Upgrade OMNI appliance

This information explains how to upgrade the OMNI appliance. You must be in the OMNI VM Console to use these steps. Once you upgrade the appliance, you must then register the appliance with the vCenter Server.

NOTE: The OMNI appliance upgrade information only applies to the OMNI 1.2 patch installation.

1. Download the OMNI upgrade image from theDell EMC Support portal, then store the image on an SCP server.
2. Select **6. Upgrade Appliance**.

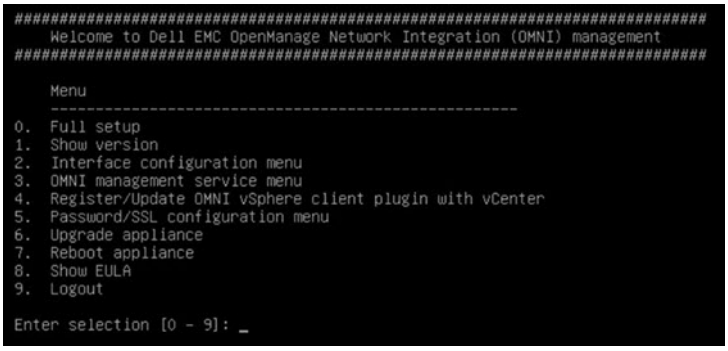

   The display lists all the applications which can be upgraded along with the old and new versions. Upgrading requires restarting the services.
3. Enter the SCP server IP/hostname, username, and the path to the upgrade .zip file and password.
4. Verify all information, then enter **Y** to continue.

```
Enter selection [0 - 10]: 7
2020-02-26 04:11:27 INFO  [setup.sh] Getting the upgrade file
Remote SCP Server IP/Hostname:
Username: admin
Path to the upgrade zip file: /home/isengard/OMNI1230.zip
The authenticity of host '                                    ' can't be established.
ECDSA key fingerprint is                                                          .
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '                ' (ECDSA) to the list of known hosts.
admin@              's password:
OMNI1230.zip    100%   10MB  30.9MB/s   00:00
                INFO [setup.sh] File Successfully copied to /home/isengard/upgrade/upgrade.zip
                INFO [setup.sh] Verifying..
Archive:  /home/isengard/upgrade/upgrade.zip
  inflating: /home/isengard/upgrade/setup.sh
  inflating: /home/isengard/upgrade/proxy.conf
  inflating: /home/isengard/upgrade/index.html
  inflating: /home/isengard/upgrade/OMNI_UserGuide.pdf
 extracting: /home/isengard/upgrade/delawareos10.zip
 extracting: /home/isengard/upgrade/version.txt
  inflating: /home/isengard/upgrade/passwd_mgr.py
  inflating: /home/isengard/upgrade/OMNI_ReleaseNotes.pdf
 extracting: /home/isengard/upgrade/sslworkspace.zip
 extracting: /home/isengard/upgrade/rls.label
 extracting: /home/isengard/upgrade/vcenterapp.zip
                INFO  [setup.sh] Nginx Proxy Conf will be upgraded
                INFO  [setup.sh] vXrail plugin file will be upgraded
                INFO  [setup.sh] OMNI Zero Touch application will be upgraded
                INFO  [setup.sh] setup File will be upgraded
                INFO  [setup.sh] New User Guide will be updated
                INFO  [setup.sh] New Release Notes will be updated
                INFO  [setup.sh] Current OMNI appliance version : 1.1.29
                INFO  [setup.sh] Current OMNI plugin version : 1.1.29
                INFO  [setup.sh] New OMNI appliance version : 1.1.29
                INFO  [setup.sh] New OMNI plugin version : 1.2.30

Upgrade will restart the service if running. Proceed? [y]?
```

```
>=3.6->vcenterapp==0.1.0)
Requirement already satisfied: more-itertools>=4.0.0 in /usr/local/lib/python3.5/dist-packages (from
 pytest<4.0,>=3.6->vcenterapp==0.1.0)
Requirement already satisfied: atomicwrites>=1.0 in /usr/local/lib/python3.5/dist-packages (from pyt
est<4.0,>=3.6->vcenterapp==0.1.0)
Requirement already satisfied: pluggy>=0.7 in /usr/local/lib/python3.5/dist-packages (from pytest<4.
0,>=3.6->vcenterapp==0.1.0)
Requirement already satisfied: pathlib2>=2.2.0; python_version < "3.6" in /usr/local/lib/python3.5/d
ist-packages (from pytest<4.0,>=3.6->vcenterapp==0.1.0)
Requirement already satisfied: setuptools in /usr/lib/python3/dist-packages (from pytest<4.0,>=3.6->
vcenterapp==0.1.0)
Requirement already satisfied: attrs>=17.4.0 in /usr/local/lib/python3.5/dist-packages (from pytest<
4.0,>=3.6->vcenterapp==0.1.0)
Requirement already satisfied: pbr>=0.11 in /usr/local/lib/python3.5/dist-packages (from mock<3.0,>=
2.0->vcenterapp==0.1.0)
Requirement already satisfied: asn1crypto>=0.21.0 in /usr/local/lib/python3.5/dist-packages (from cr
yptography>=2.2.1->pyopenssl<19.0.0,>=18.0.0->vcenterapp==0.1.0)
Requirement already satisfied: cffi!=1.11.3,>=1.7 in /usr/local/lib/python3.5/dist-packages (from cr
yptography>=2.2.1->pyopenssl<19.0.0,>=18.0.0->vcenterapp==0.1.0)
Requirement already satisfied: MarkupSafe>=0.23 in /usr/local/lib/python3.5/dist-packages (from Jinj
a2>=2.10->flask<2.0,>=1.0->vcenterapp==0.1.0)
Requirement already satisfied: pycparser in /usr/local/lib/python3.5/dist-packages (from cffi!=1.11.
3,>=1.7->cryptography>=2.2.1->pyopenssl<19.0.0,>=18.0.0->vcenterapp==0.1.0)
Installing collected packages: vcenterapp
Successfully installed vcenterapp-0.1.0
'setup.sh' -> '/home/isengard/setup.sh'
'rls.label' -> '/home/isengard/rls.label'
2018-10-19 12:42:15 INFO  [setup.sh] Starting all services
2018-10-19 12:42:15 INFO  [setup.sh] OMNI Zero Touch Management Application Service Start successful
2018-10-19 12:42:15 INFO  [setup.sh] OMNI Application Server Service Start successful
2018-10-19 12:42:15 INFO  [setup.sh] OMNI Web Server Service Start successful
2018-10-19 12:42:15 INFO  [setup.sh] Removing upgrade files
2018-10-19 12:42:16 INFO  [setup.sh] Plugin need to be update registration to vCenter for new UI cha
nges
2018-10-19 12:42:16 INFO  [setup.sh] Session will be closed now. Please log back in.
press [enter] to continue...
```

5. Select **4. Register/Update OMNI vSphere client plugin with vCenter** to register the plug-in.
6. Enter the FQDN to use for registration, then repeat the steps to update the plug-in with the vCenter Server.

Upgrade from OMNI 1.1 to OMNI 1.2

1. Unregister the OMNI 1.1 plug-in using the OMNI CLI option.
2. Shut down the OMNI 1.1 VM.
3. Deploy the OMNI 1.2 OVA, then register the OMNI plug-in with the vCenter (see OMNI vCenter client plug-in registration).

4. (Optional) Delete the OMNI 1.1 VM, if needed.

Upgrade SmartFabric OS in switch

This information describes how to upload an image to upgrade the fabric. For information about changing the SmartFabric password, upgrading the OS or replacing the switch, see Summary.

NOTE: Fabric upgrade must be performed during OMNI maintenance mode (see OMNI maintenance mode).

Upload image

1. Select **Service Instance > Summary > Fabric Actions > Upgrade OS**.

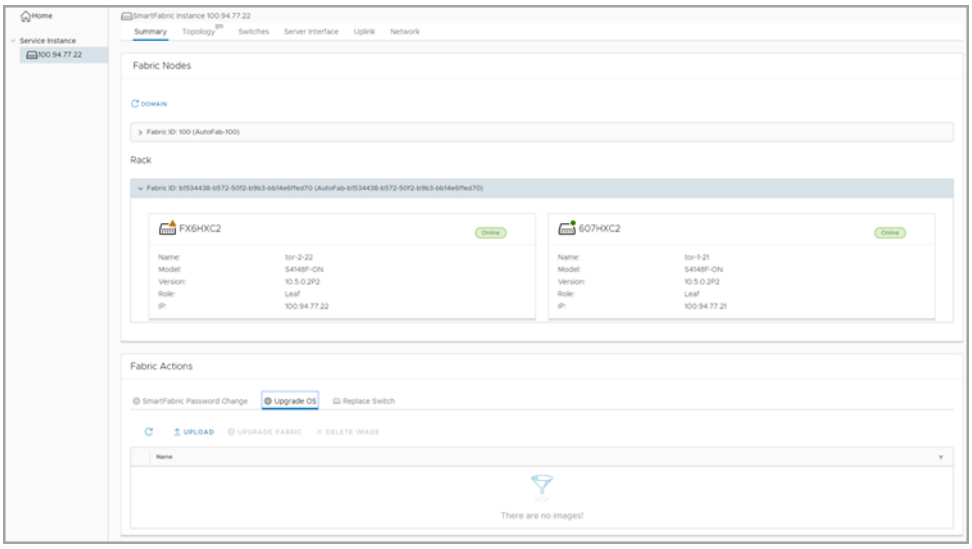

2. Click **Upload** to upload the .bin file.

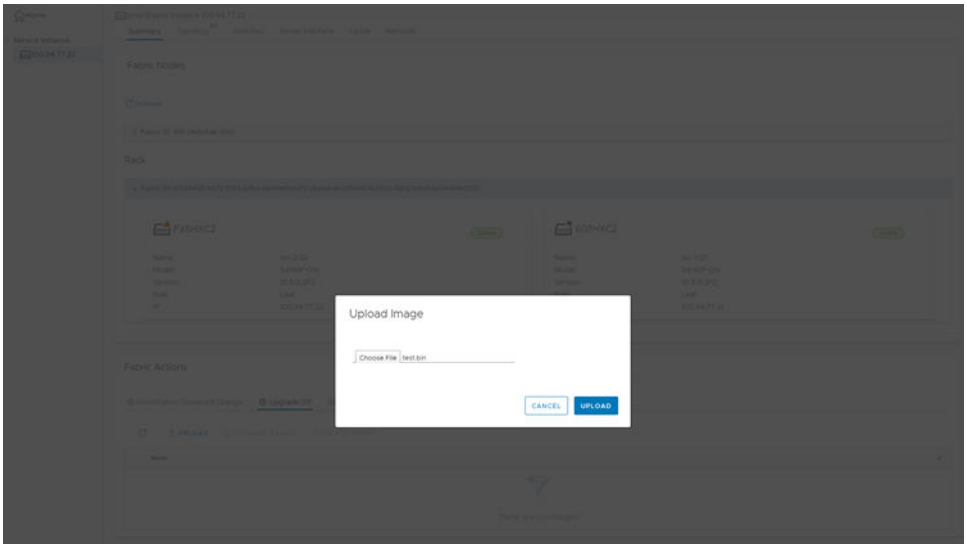

Upgrade fabric

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

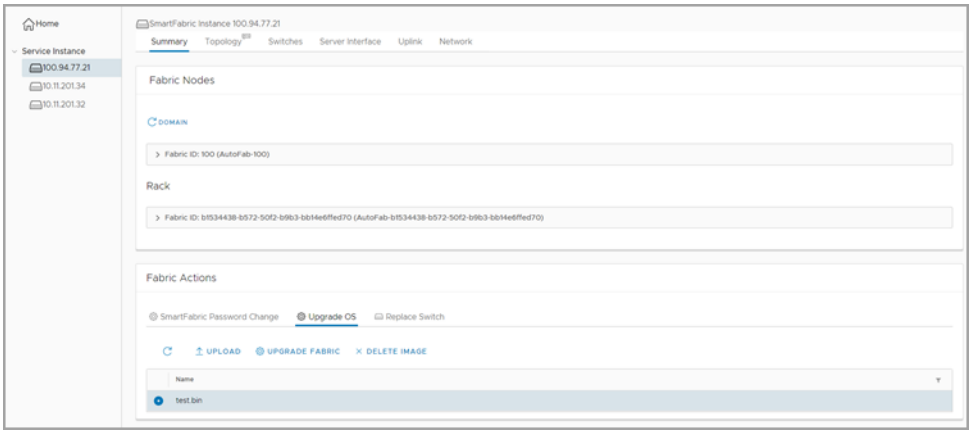

NOTE: Upgrade Fabric option upgrades all the switches in a network fabric. You cannot stop the upgrade after it is triggered.

2. Confirm the upgrade; click **Upgrade**.

NOTE: Maintenance mode is a precautionary step during the SmartFabric upgrade. If you continue, the upgrade stops OMNI automatic VLAN configuration.

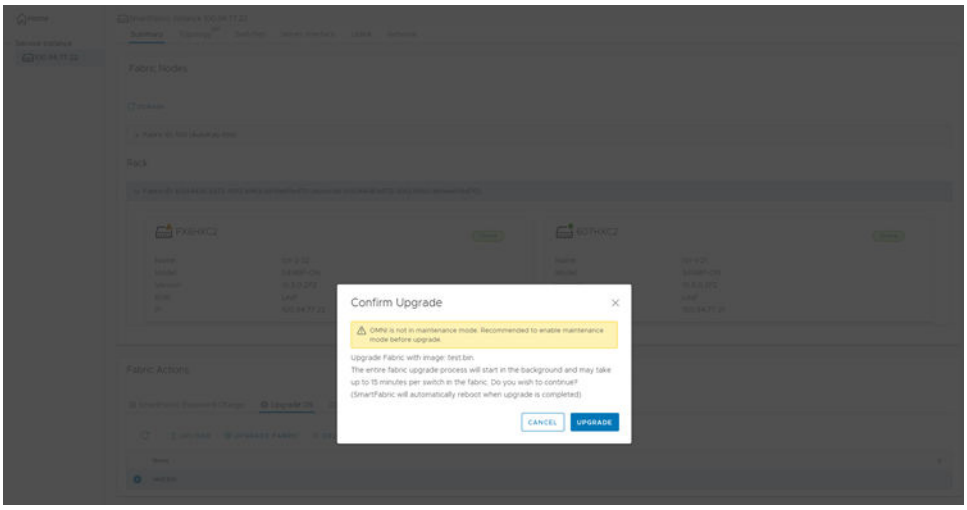

3. Fabric upgrade success; close the message. SmartFabric automatically reboots when the upgrade is complete.

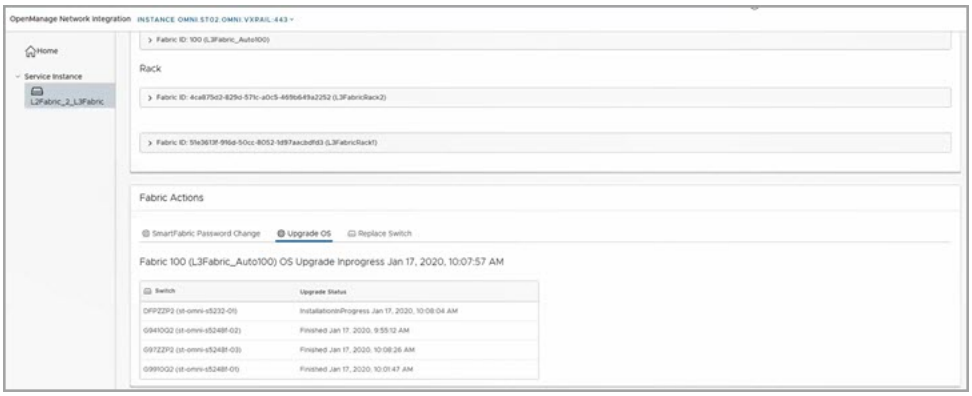

NOTE: The system displays the status of the network fabric upgrade only during the upgrade, and the status disappears from the screen after the upgrade is completed.

Delete image

1. Select the .bin image to delete.

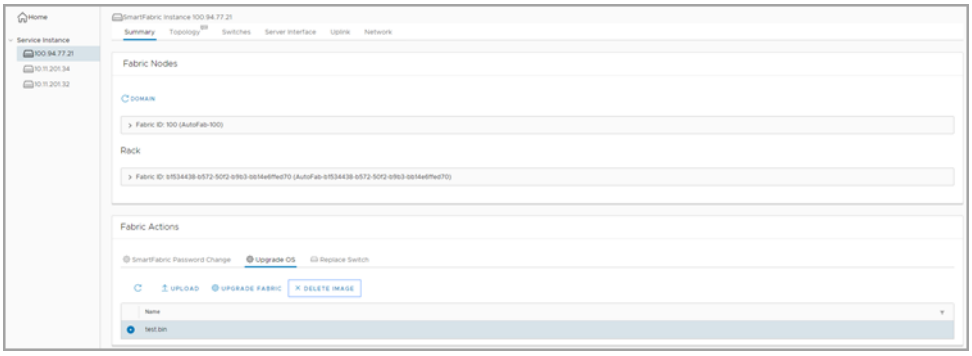

2. Click **Delete Image**, then click **Delete**.

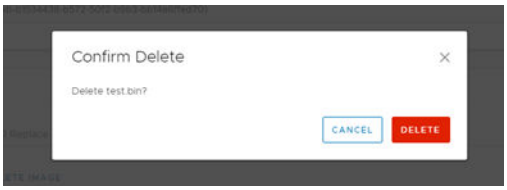

3. Once the image is deleted, a success message displays; close the window.

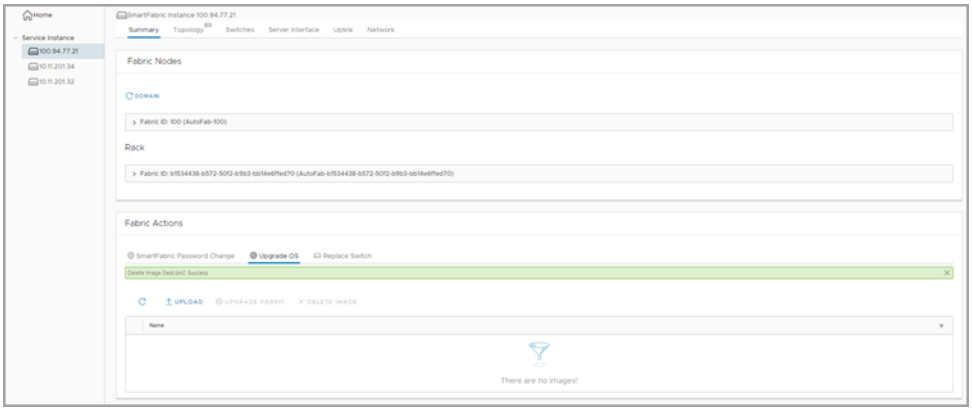

Replace switch in a fabric

You can replace an faulty OS10 switch in a SmartFabric environment.

Follow the below instructions 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. Take a backup of the following configurations from the old 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. The new switch must have the same SmartFabric OS10 version as the existing 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. You can use the `smartfabric l3fabric enable role` command to enable SmartFabric Services. 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, go to **OMNI** > **Summary** > **Replace Switch** to clean up the old switch and complete the switch replacement workflow.

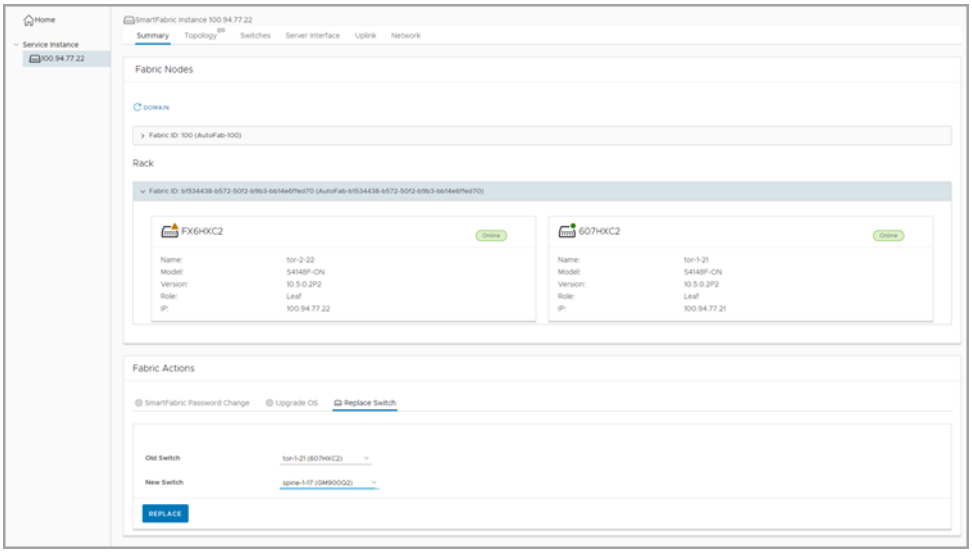

11. Select the switch to replace from the list, select the new switch, and click **REPLACE**. The system displays replace success message.

9

OMNI support

This information describes how to access the OMNI support. Select **Support** to view the support options.

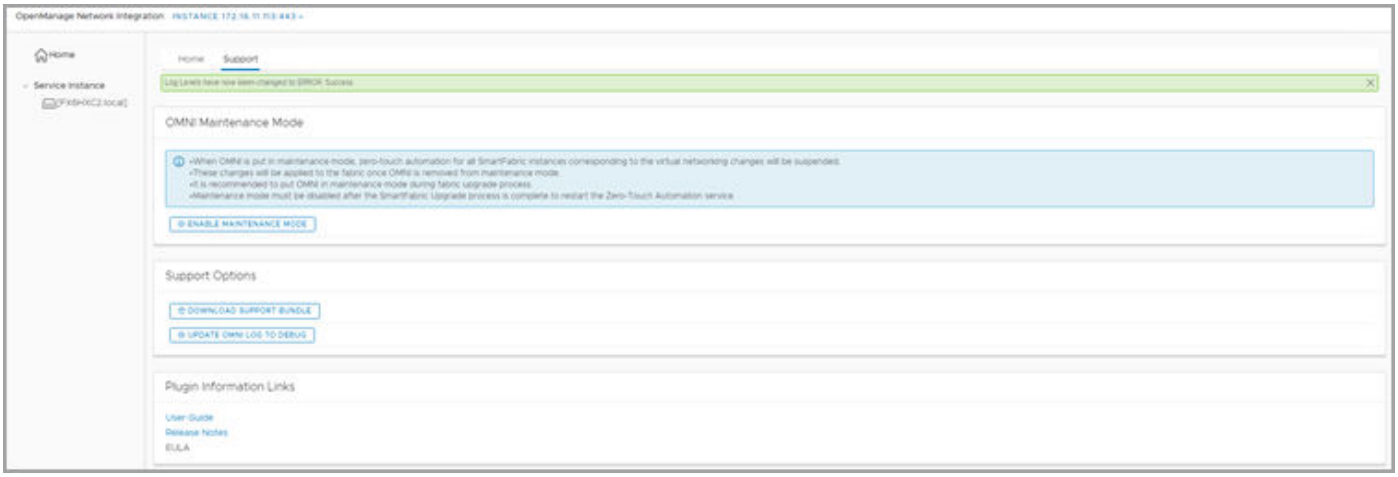

OMNI maintenance mode

Enable Maintenance Mode

NOTE: It is recommended to enable Maintenance Mode before upgrading the network fabric.

Enabling Maintenance Mode disables zero-touch automation for all SmartFabric instances. Enabling Maintenance Mode prevents OMNI from configuring networks on SmartFabrics when there are changes in the vCenter port groups.

1. Click **Support > Enable Maintenance Mode**.

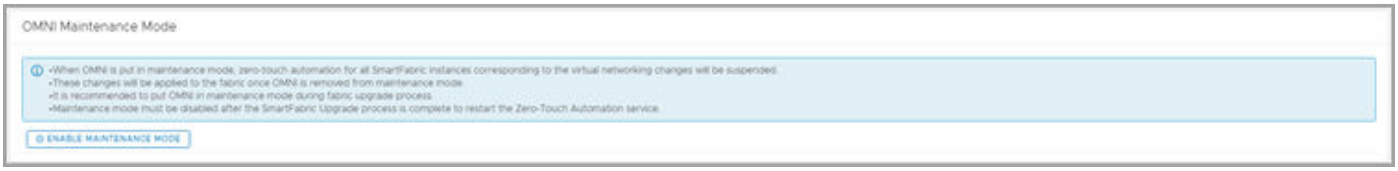

2. Click **Enable Maintenance Mode** to confirm.

3. Maintenance Mode enabled; close the message.

Disable Maintenance Mode

Disabling Maintenance Mode restarts zero-touch automation for all SmartFabric instances. OMNI begins to create networks on the SmartFabric instances based on the port groups in the vCenter.

1. Click **Support > Disable Maintenance Mode**.
2. Click **Disable Maintenance Mode** to confirm.

3. Maintenance Mode disabled; close the message.

Support options

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

When OMNI's log-level is set to ERROR only, error logs are recorded. When OMNI's log-level is set to DEBUG, error logs and logs with additional information is recorded. The DEBUG level should only be selected while trying to diagnose an issue.

1. Click **Download Support Bundle**.

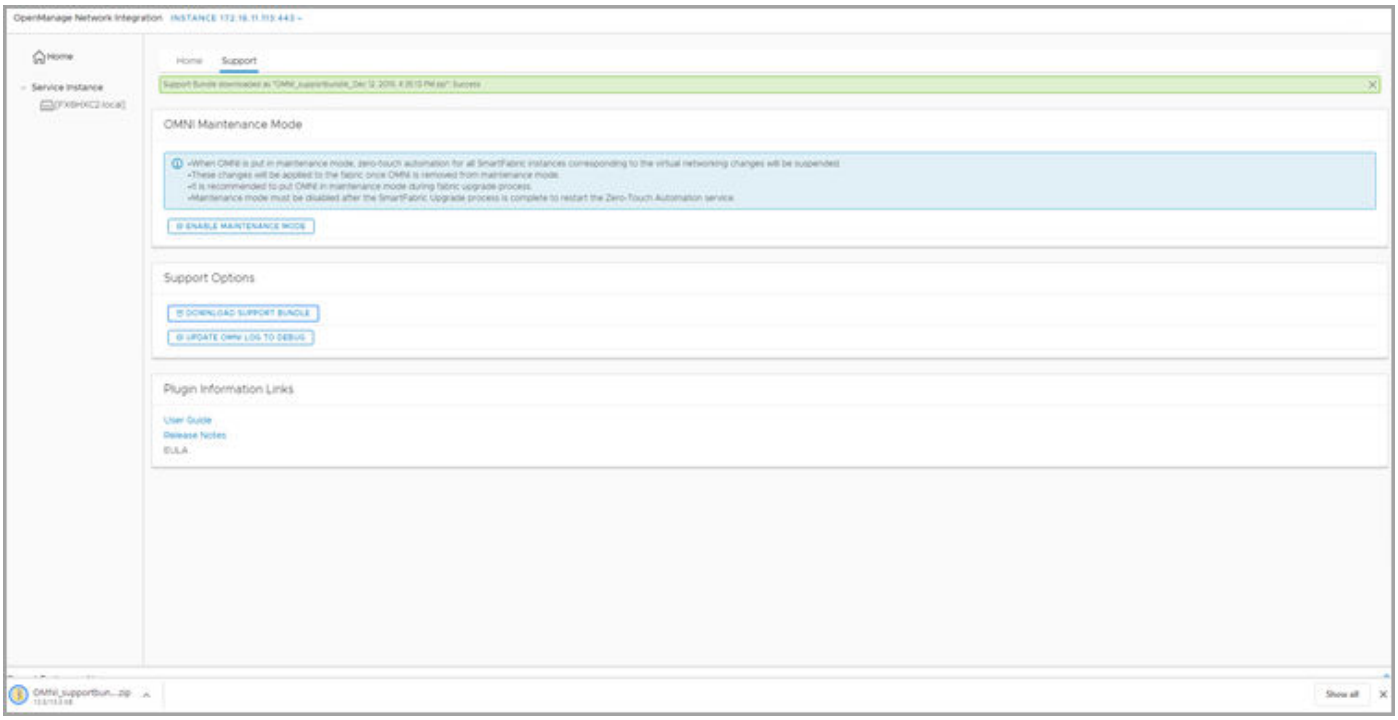

2. (Optional) Click **Update OMNI Log To Debug** to modify the log-level to Error.

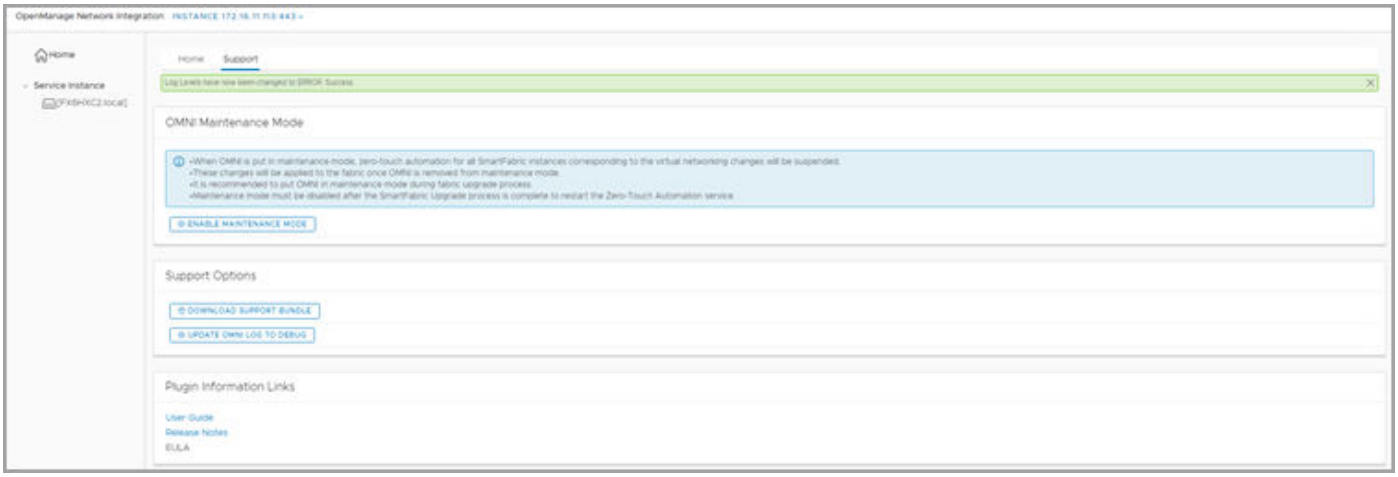

3. (Optional) Click **Update OMNI Log To Error** to modify the log-level to Debug.

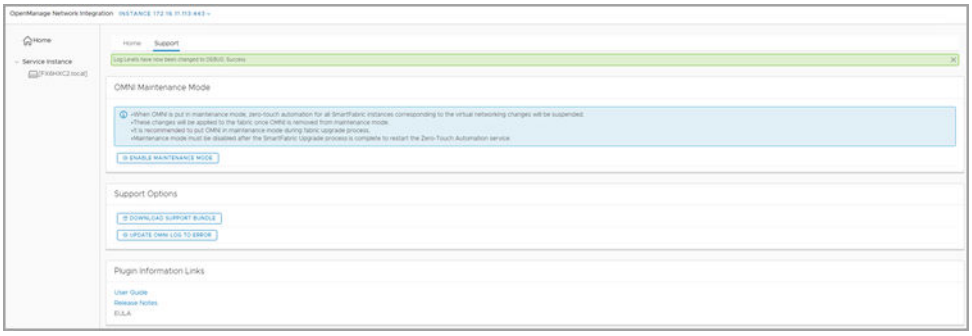

Plugin information links

You can view the User Guide, Release Notes, or end-user license agreement (EULA) to view or download.

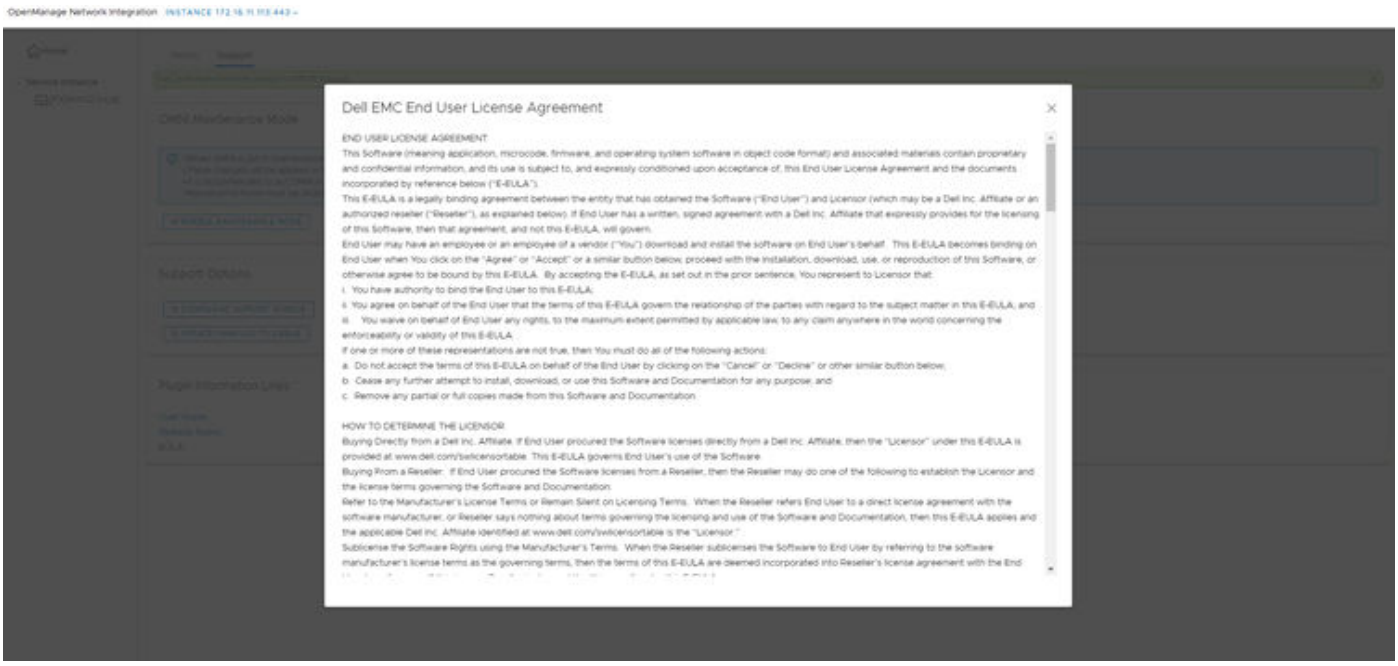

10

Troubleshooting

You can use the following information to troubleshoot the SmartFabric vCenter appliance connectivity, SmartFabric errors, and UI population errors.

SmartFabric appliance connectivity

Check the IP and DNS settings and connection status.

1. Select **2. Interface Configuration Menu**.

```
################################################################################
    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.  Register/Update OMNI vSphere client plugin with vCenter
5.  Password/SSL configuration menu
6.  Upgrade appliance
7.  Reboot appliance
8.  Show EULA
9.  Logout

Enter selection [0 - 9]: 2
```

2. Select **1. Show Interfaces** (q to close view).

```
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
sudo: unable to resolve host omni: No such file or directory
ens160: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 172.16.11.113  netmask 255.255.255.0  broadcast 172.16.11.255
        inet6 fe80::250:56ff:febc:afa3  prefixlen 64  scopeid 0x20<link>
        ether 00:50:56:bc:af:a3  txqueuelen 1000  (Ethernet)
        RX packets 1386077  bytes 683505104 (651.8 MiB)
        RX errors 0  dropped 47  overruns 0  frame 0
        TX packets 188512  bytes 201083390 (191.7 MiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

ens192: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet6 fe80::598:2e31:8941:cd4e  prefixlen 64  scopeid 0x20<link>
        ether 00:50:56:bc:9d:76  txqueuelen 1000  (Ethernet)
        RX packets 125203  bytes 84245483 (80.3 MiB)
        RX errors 0  dropped 24  overruns 0  frame 0
        TX packets 59624  bytes 10264094 (9.7 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 1  (Local Loopback)
        RX packets 107732  bytes 497910397 (474.8 MiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 107732  bytes 497910397 (474.8 MiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

3. Select **2. Show Connection Status**.

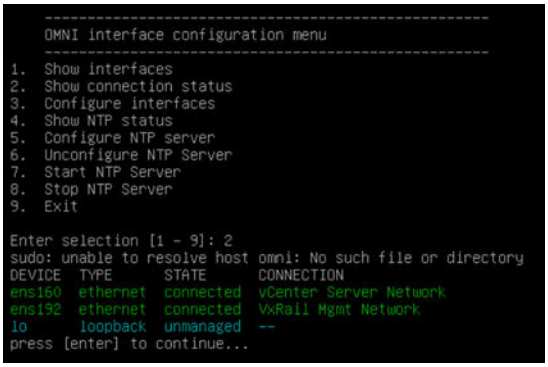

4. Issue a curl https request from the SmartFabric vCenter appliance console to the plug-in appliance IP configured on ens160.

SmartFabric error

Check ens192 has IPv6 link-local address that is enabled and is up.

NOTE: Multiple fabric or vCenter registrations are shown in Status SmartFabric Management Service.

1. Select **2. Status OMNI Managment Service**.

   All services must be active. If they are not all active, restart the service by selecting **4. Restart OMNI Management Service**. Confirm that the error is resolved.
2. Confirm the password that is supplied during registration is correct. If not correct, re-register the plug-in to update the SmartFabric password on the appliance.

Configure NTP server

NOTE: If the NTP server is not configured, the SmartFabric vCenter appliance VM will not synchronize with the data center.

1. Select **2. Interface Configuration Menu**.

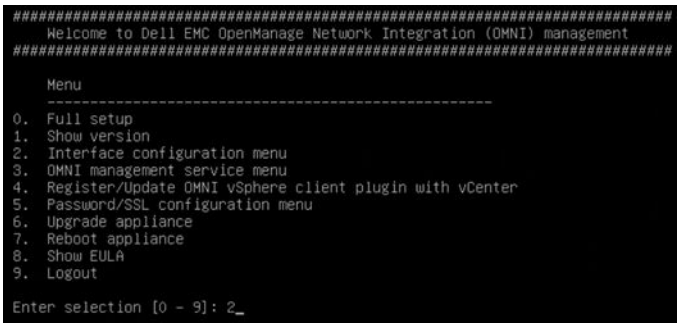

2. Select **5. Configure NTP Server**, then enter the **NTP Server IP/Hostname**.

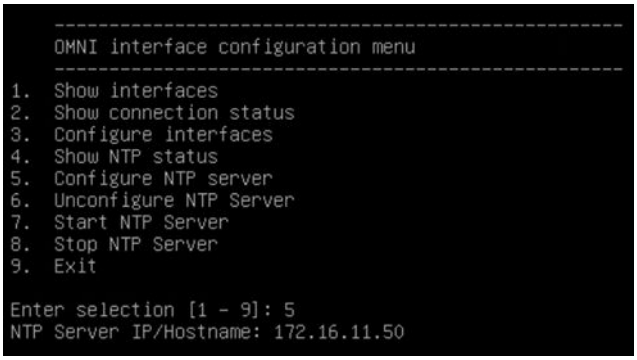

NTP server status

1. Select **2. Interface Configuration Menu**.

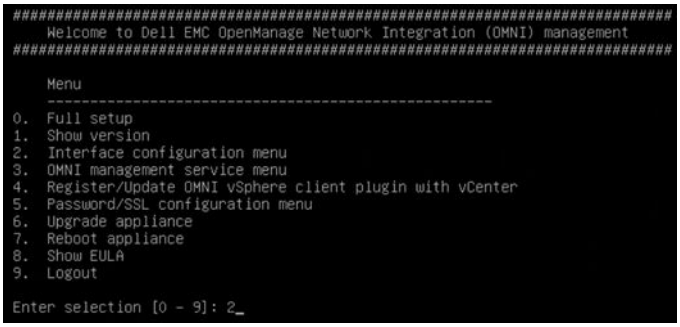

2. Select **4. Show NTP Status**.

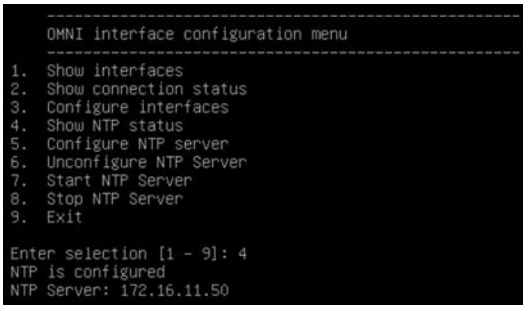

3. Enter the **Validate NTP Server IP address/host name**.

UI is not populated

NOTE: If any IP address or SSL certificate changes on the SmartFabric VM, OMNI automation services can be restarted by enabling Maintenance mode, then disabling Maintenance mode (see OMNI support).

Check the service status on the plug-in VM.

1. Select **4. Restart SmartFabric Management Service**.
2. Select **2. Status SmartFabric Management Service** to list the registered vCenter and SmartFabric OMNI the VM is managing. Confirm that all services are active.

Create support bundle

NOTE: We recommend using OpenManage Network Integration > Home > Support > Update OMNI Log to Debug option.

1. Select **4. SmartFabric Management Service Menu**.

2. Select **5. Create Support Bundle** to create a support bundle at /tmp/support-bundle.tar.gz on the SmartFabric OMNI VM.

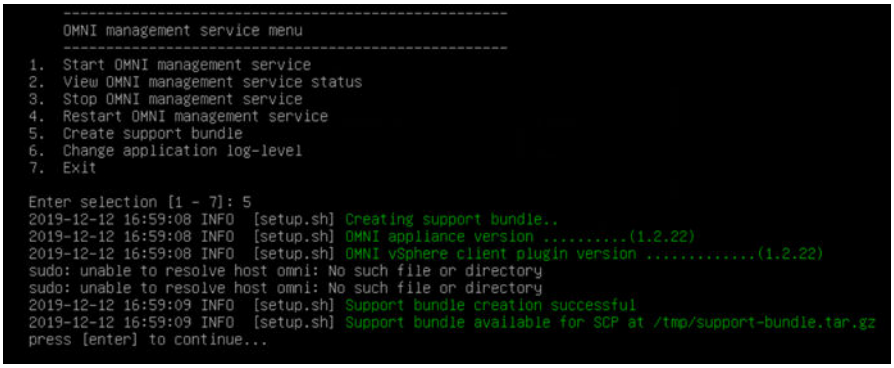

3. From an external host, scp using *admin* to transfer the support bundle file out. SCP credentials for the SmartFabric appliance are the same as the SmartFabric appliance console password. By default, the username and password is admin/admin.

NOTE: Recommendation is to set the log level change to DEBUG before creating the support bundle.

Change log level

NOTE: We recommend using OpenManage Network Integration > Home > Support > Update OMNI Log to Debug option.

1. Select **4. SmartFabric Management Service menu**.
2. Select **6. Change Application Log Level** to display the current log level and switch accordingly.

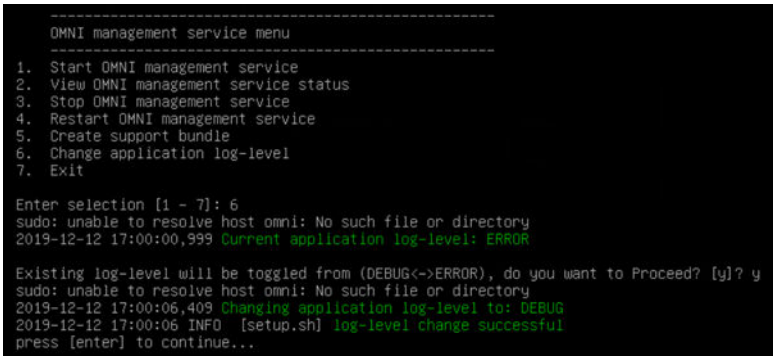

NOTE: By default, the SmartFabric appliance comes with a log-level of ERROR. The appliance log can be swapped between ERROR to DEBUG, or the opposite way.

3. Stop if the log level is already on the wanted log level.

Reset SmartFabric VM password

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

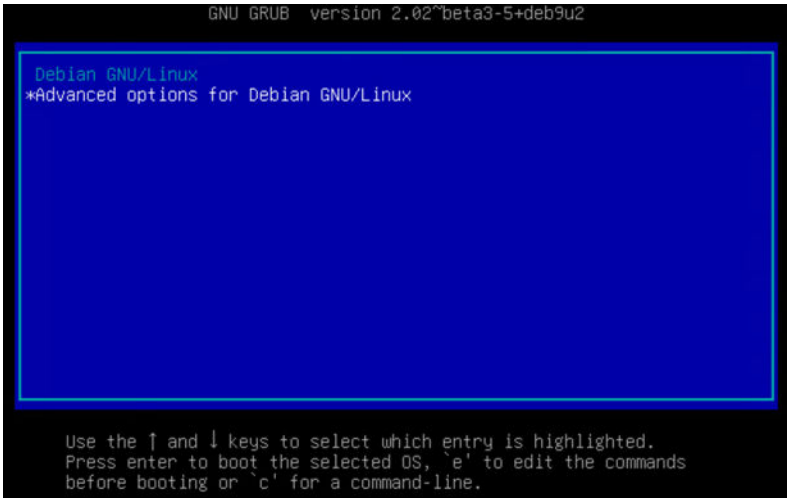

2. Use the arrow keys to go to the line starting with `linux` and ending with `ro quiet`.

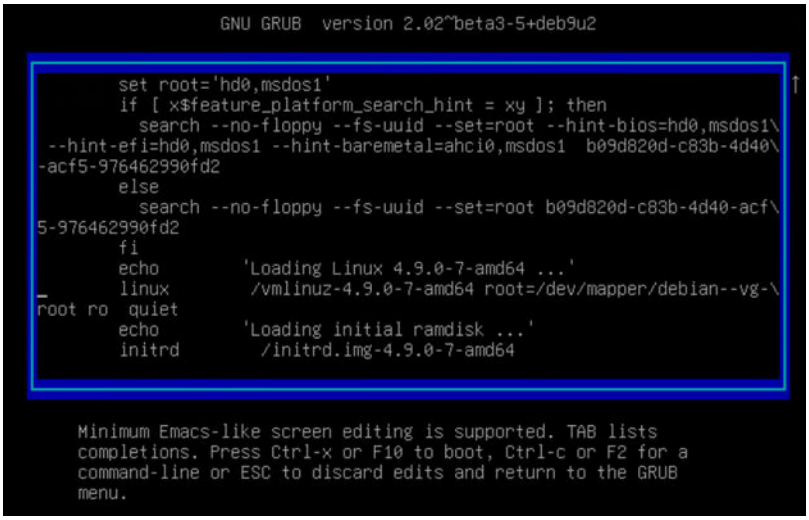

3. Append `init=bin/bash` after `ro quiet`.

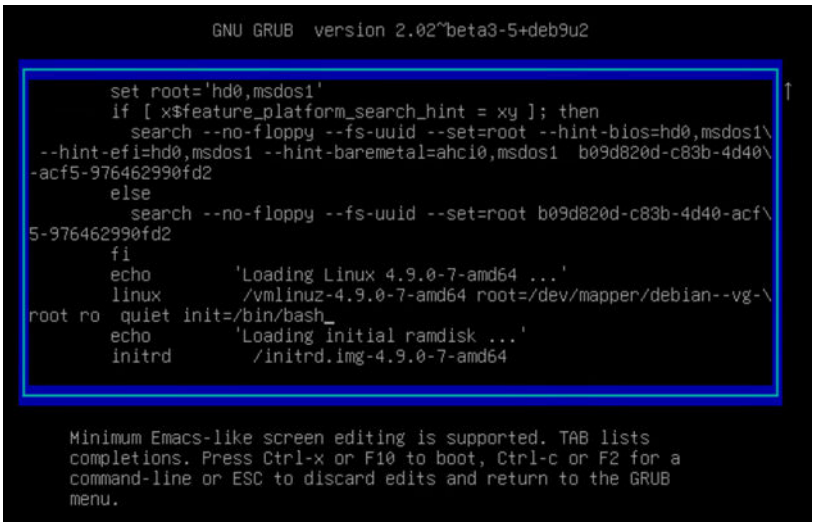

4. Press **Ctrl-X** to boot into the shell with root access.

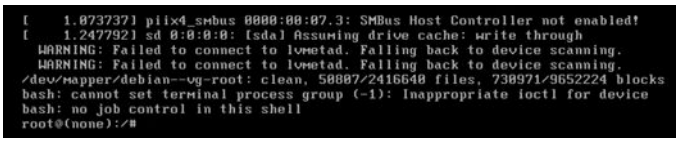

5. Remount the / directory.

```
# mount / -re -o remount
```

6. Change the password for admin using `passwd admin`, then enter and confirm the new password.

```
root@(none):/# mount / -rw -o remount
root@(none):/# passwd admin
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
root@(none):/#
```

7. Reset the VM from vCenter. You could log in through the new password for the SmartFabric VM.

Missing networks on server interfaces

If OMNI fails to create and associate the appropriate network on a server interface during automation, OMNI automation services can be restarted so that OMNI reconfigures the networks.

OMNI automation services can be restarted by enabling Maintenance mode, then disabling Maintenance mode (see OMNI support).

OMNI unable to resolve vCenter FQDN

A change in the DNS can cause an issue during FQDN resolution. Any change in DNS must be updated on the application through Option 2. Interface configuration Menu. You must then set the proper DNS for the interface. For complete information, see *Network interface profile configuration* in OpenManage Network Integration.

Certificate not trusted error

If OMNI is having issues communicating with the vCenter due to SSL certificate errors, new SSL certificates must be installed.

1. To install new SSL certificates, see *OMNI vCenter client plugin-in registration* in OpenManage Network Integration.
2. OMNI automation services can be restarted by enabling Maintenance mode, then disabling Maintenance mode (see OMNI support).

Index

P

R

S

T

U

V