Dell EMC SmartFabric Director

User Guide Release 1.2

May 2020
Rev. A02

Notes, cautions, and warnings

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

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

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

Contents

1

Revision history

This table provides an overview of the changes in this guide.

Table 1. Revision history

Release	Revision	Description
1.1.0	A00 (January 2020)	Initial release
1.1.1	A01 (March 2020)	• Define switch lifecycle job — Enter a user-specified OS10 image name • Upgrade SFD — Upgrade SFD software
1.2.0	A02 (May 2020)	• Installation using vCenter 6.7 — Updated information on service tag input • Specify system settings — Added new section to upload the service tag • Create user accounts — Updated screenshots for user management • Define fabric intent — Added BFD to Layer 3 configuration

2

SmartFabric Director

Dell EMC SmartFabric Director (SFD) enables data center operators to build, operate, and monitor an open network underlay fabric. SFD works with Dell EMC PowerSwitch Series switches to ensure that their physical underlay networks are tuned for the specific overlay environment.

SFD enables the physical switch underlay infrastructure to keep pace with the changing demands of virtualized and software-defined networks, and provides customers a single view for operating, managing, and troubleshooting of physical and virtual networks.

Features

- Abstracted view of the fabric—no must manage individual switches
- Define, build, and maintain a Layer 2 or Layer 3 leaf spine data center fabric (underlay)
- Intent template-based provisioning underlay
- Authoritative repository of intent and switch configuration and state
- Fabric health management and monitoring including events, logs, alarms, states, and metrics (counters)
- Operator-driven remediation
- Full life-cycle management of switches including grouping of switches and scheduling of jobs
- Uses OpenConfig (gNMI, gNOI) for provisioning and streaming telemetry of switches

Inputs

- Provisioning using REST or gRPC/gNOI
- ONIE and gNOI life-cycle management
- Streaming telemetry using gRPC
- Agentless or Agent interface to switches
- L2 or L3 fabric topology

Streaming telemetry

Model-driven telemetry is a new approach for network monitoring. Data is streamed from network switches continuously, using a push model which provides near real-time access to operational statistics. Applications can subscribe to specific data items they need, by using standard-based YANG data models.

Streaming telemetry enables users to push data off the switch to an external collector at a higher frequency, more efficiently, and data on-change streaming.

Models

- *destination-group* tells the switch where to send telemetry data and how
- *sensor-group* identifies a list of YANG models that the switch should stream
- *subscription-profile* ties together the destination-group and the sensor-group

3

Getting started

This information describes the component and configuration requirements.

Dell EMC SmartFabric Director

- Dell EMC SmartFabric Director release 1.2.0

SmartFabric OS10

All PowerSwitches must be running Dell EMC SmartFabric OS10:

Release 1.2.0	10.5.1.2
Release 1.1.2	10.5.0.4 or 10.5.0.5
Release 1.1.1	10.5.0.4 or 10.5.0.5
Release 1.1.0	10.5.0.4

Dell EMC PowerSwitches

- S4048-ON, S4048T-ON
- S4112F-ON, S4112T-ON
- S4128F-ON, S4128T-ON
- S4148F-ON, S4148FE-ON, S4148T-ON
- S4248FB-ON, S4248FBL-ON
- S5212F-ON
- S5224F-ON
- S5232F-ON
- S5248F-ON
- S5296F-ON
- S6010-ON
- Z9100-ON
- Z9264F-ON

VMware requirements

VMware ESXI

- Virtualization-ready x86 server
- VMware ESXi 6.7 U1, U2 (recommended); ESXi 6.5, U1, U2, U3
- VMware vSphere Enterprise Plus license
- Virtual appliance (OVA)
- 4vCPU
- 16G memory
- 100G available disk space (higher disk sizes may be required depending on fabric size and data retention requirements)

VMware NSX-T

See docs.vmware.com/VMware NSX-T Data Center for complete NSX-T requirements.

More requirements

- Web browser — Chrome (version 72.0.3626.121 and later) and Firefox (version 68.0 and later) recommended
- vSphere web client 6.5 U1, U2, U3 — supported for Flash client; not supported for HTML5 client
- vSphere web client 6.7 (all versions) — nonsupported for Flash client; supported for HTML5 client
- Text or JSON editor to modify the JSON wiring diagram if required

4

Software installation

This information describes how to install SFD in your SmartFabric OS10 network. If your switch came preinstalled with Dell EMC SmartFabric OS10, see Log in to SmartFabric OS10.

NOTE: For detailed hardware installation steps, see the product-specific *Installation Guide* at www.dell.com/support/.

Topics:

- Download SFD image
- Log in to SmartFabric OS10
- Upgrade SFD

Download SFD image

This information explains show to download the SmartFabric Director software image.

1. Sign into DDL using your account credentials.
2. Locate your entitlement ID and order number, then select the product name.
3. Select the **Products** tab and view your service tag that is located under Associated Hardware of Software ID; write the service tag down. The service tag is needed during first-time setup. The service tag is also visible in the license key.
4. Select the **Available Downloads** tab, select the wanted files to download, then click **Download**.
5. Read the Dell End-User License Agreement. Scroll to the end of the agreement, then click **Yes, I agree**.
6. Select how to download the software files, then click **Download Now**.
7. After you download the image, unpack the .tar file on a Linux or Windows server, then open the README file for instructions on how to validate the OVA file.

NOTE: The available downloads include the software image, release notes, user guide, and JSON wiring diagram template (see Fabric wiring diagram definition for complete information).

Log in to SmartFabric OS10

To log in to SmartFabric OS10, turn on the device and wait for the system to perform a power-on self-test (POST). Enter `admin` for both the default username and user password.

For better security, change the default `admin` password during the first SmartFabric OS10 login. The system saves the new password for future logins. After you change the password through the CLI, enter the `write memory` command to save the configuration.

```
OS10 login: admin
Password: admin
Last login: Thu Dec 12 13:58:27 2019 on ttyS0

The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.


-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-
-*          Dell EMC Network Operating System (OS10)            *-
-*                                                               *-
-* Copyright (c) 1999-2020 by Dell Inc. All Rights Reserved.    *-
-*                                                               *-
-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-*-

This product is protected by U.S. and international copyright and
intellectual property laws. Dell EMC and the Dell EMC logo are
```

```
trademarks of Dell Inc. in the United States and/or other
jurisdictions. All other marks and names mentioned herein may be
trademarks of their respective companies.

OS10# write memory
```

Check SmartFabric OS10 version

1. View the SmartFabric OS10 version in EXEC mode.

```
OS10# show version
Dell EMC Networking OS10-Enterprise
Copyright (c) 1999-2020 by Dell Inc. All Rights Reserved.
OS Version: 10.5.0.4
Build Version: 10.5.0.4.433
Build Time: 2019-12-12T22:18:40-0700
System Type: S4148F-ON
Architecture: x86_64
Up Time: 2 days 03:37:25
```

2. (Optional) If your switch is not preloaded with SmartFabric OS10 10.5.0.4, you must upgrade the operating system (see Upgrade SmartFabric OS10 in the *Dell EMC SmartFabric OS10 User Guide*).

Upgrade SFD

This information explains how to upgrade SFD to the next software release.

NOTE: If you are running SFD 1.1.0 or 1.1.2, you can directly upgrade to 1.2.0. If you are running SFD 1.1.1, you must first upgrade to 1.1.2, then upgrade to 1.2.0. The SFD software cannot be downgraded to a lower version.

1. Download the upgrade bundle, then store it locally on SFD, or in a remote location (see Download SFD image).

   NOTE: The SFD image name and upgrade bundle name should not be modified. The installation or upgrade will fail if you change the filenames.

2. SSH to the IP address configured for SFD.

```
login-srv-05-user%:~> ssh username@sfd.local@ip_address
admin@sfd.local@10.12.124.125's password:
Last login: Wed May  6 18:00:59 2020 from 10.12.1.9
```

3. Log in to the SFD CLI (see Using the CLI).

```
admin@sfd.local@SFD-R5:~$ sfd
DellEMC SmartFabric Director CLI
sfd>
```

4. Upgrade the SFD software from an absolute path on a local server.

```
sfd> upgrade absolute path of upgrade bundle
```

   or

   Upgrade the SFD software from a remote host.

```
sfd> upgrade absolute path of upgrade bundle on remote host --server remote server IP --
username username --password password
```

   On successful upgrade, `Upgrade successful!` displays.

5. Verify the upgrade, and compare the current installed software version with the intended SFD version.

```
sfd> cat /opt/vmware/nfc/version/NFC_VERSION
```

For complete information about the `upgrade` command, see upgrade.

5

Switch configuration

This information explains how to configure SmartFabric OS10 before installing SFD. For complete configuration information, see the *Dell EMC SmartFabric OS10 User Guide*.

Topics:

Management interface

This information explains how to configure Management interface access to network devices. You can configure the Management interface, but the configuration options on this interface are limited. You cannot configure gateway addresses and IP addresses if it appears in the main routing table. Proxy ARP is not supported on this interface.

1. Configure the Management interface in CONFIGURATION mode.

```
interface mgmt 1/1/1
```

2. By default, DHCP client is enabled on the Management interface. Disable DHCP client operations in INTERFACE mode.

```
no ip address dhcp
```

3. Configure an IP address and mask on the Management interface in INTERFACE mode.

```
ip address A.B.C.D/prefix-length
```

4. Enable the Management interface in INTERFACE mode.

```
no shutdown
```

Configure Management interface

```
OS10(config)# interface mgmt 1/1/1
OS10(conf-if-ma-1/1/1)# no ip address dhcp
OS10(conf-if-ma-1/1/1)# ip address 10.1.1.10/24
OS10(conf-if-ma-1/1/1)# no shutdown
```

For complete information about configuring Management interfaces, see the *Dell EMC SmartFabric OS10 User Guide Release 10.5.0* at www.dell.com/support/.

Crypto security

This information explains how to prepare your switch for SmartFabric Director from the OS10 side to install the crypto security profile and license.

The gNMI agent, available with SmartFabric OS10 release 10.5.0.1 and later, provides a new interface to configure OS10 devices. It uses gNMI protocol and OpenConfig YANG models to support create, read, update, and delete (CRUD) operations, life cycle management through gNOI and configuration of streaming telemetry.

The gNMI agent listens to SFD to receive remote configuration-change requests or upgrade and downgrade instructions. As a part of these remote configuration changes, the gNMI agent enables the telemetry agent to transmit preconfigured sensor group data in the OpenConfig format to SFD.

Setup crypto security

1. Log in to SmartFabric OS10, then verify that the installed software version meets the requirements.

```
OS10# show version
Dell EMC Networking OS10 Enterprise
Copyright (c) 1999-2020 by Dell Inc. All Rights Reserved.
OS Version: 10.5.0.4
Build Version: 10.5.0.4.433
Build Time: 2019-09-26T03:50:26+0000
System Type: S4148F-ON
Architecture: x86_64
Up Time: 00:07:08
```

2. Verify your switch operating mode.

```
OS10# show switch-operating-mode

Switch-Operating-Mode : Full Switch Mode
```

3. Set up the crypto security profile and certificate, then replace `gnmi-os10-0` with a security-profile name of your choice; format is `gnmi-xxx-0` where `xxx` is any string.

```
OS10(config)# crypto security-profile gnmi-os10-0
OS10(conf-sec-profile)# certificate gnmi-os10-0
OS10(conf-sec-profile)# exit
```

4. Start `restconf`.

```
OS10(config)# rest api restconf
```

5. Set up the `gnmi-security-profile`.

```
OS10(config)# gnmi-security-profile gnmi-os10-0
OS10(config)# exit
```

6. Create crypto certificate.

```
OS10# crypto cert generate self-signed cert-file home://gnmi-os10-0.crt key-file home://
gnmi-os10-0.key cname os10
Processing certificate ...
Successfully created certificate file and key
```

7. Install the certificate.

```
OS10# crypto cert install cert-file home://gnmi-os10-0.crt key-file home://gnmi-os10-0.key
Processing certificate ...
Certificate and keys were successfully installed as "gnmi-os10-0.crt" that may be used in
a security profile. CN = os10
```

8. Set the switch operating mode to SFD, then verify the mode.

```
OS10# configure terminal
OS10# switch-operating-mode sfd
OS10# exit
OS10# show switch-operating-mode
Switch-Operating-Mode : SmartFabric Director Mode
OS10# show sfd status
Controller IP                           Port            Status
--------------------------------------------------------------------------------
::ffff:100.94.25.246                    8443            active
```

9. Save the configuration.

```
OS10# write mem
```

10. SFD displays an informational message to reload the device for SFD mode to take effect; reload the switch.

```
OS10# reload
```

For complete information about crypto security profiles, see the *Dell EMC SmartFabric OS10 User Guide Release 10.5.0*.

Switch-port profiles

This information explains switch-port profiles. A port profile determines the enabled front-panel ports and supported breakout modes on Ethernet and unified ports. Change the port profile on a switch to customize uplink and unified port operation, and the availability of front-panel data ports.

To change the port profile at the next reboot, use the `switch-port-profile` command with the wanted profile, save it to the startup configuration, then use the `reload` command to apply the changes.

1. Configure a platform-specific port profile in CONFIGURATION mode. For a stand-alone switch, enter `1/1` for `node/unit`.

   NOTE: Switch-port profiles are platform-specific. If `switch-port-profile` is not available, the configuration is not available for your specific platform.

   ```
   switch-port-profile node/unit profile
   ```

2. Save the port profile change to the startup configuration in EXEC mode.

   ```
   write memory
   ```

3. Reload the switch in EXEC mode.

   ```
   reload
   ```

The switch reboots with the new port configuration and resets the system defaults, except for the switch-port profile and these configured settings:

- Management interface 1/1/1 configuration
- Management IPv4/IPv6 static routes
- System hostname
- Unified forwarding table (UFT) mode
- ECMP maximum paths

You must manually reconfigure other settings on the switch after you apply a new port profile and reload the switch.

NOTE: After you change the switch-port profile, do not immediately back up and restore the startup file. You must use the `write memory` command and reloading the switch using the `reload` command or the new profile does not take effect.

Configure port profile

```
OS10(config)# switch-port-profile 1/1 profile-6
OS10(config)# exit
OS10# write memory
OS10# reload
```

Verify port profile

```
OS10(config)# show switch-port-profile 1/1

|  Node/Unit  |  Current    |  Next-boot  | Default   |
|-------------+-------------------+-------------------|
|     1/1     |  profile-2  |  profile-2  | profile-1 |

Supported Profiles:
profile-1
profile-2
profile-3
profile-4
profile-5
profile-6
```

For complete information about configuring specific ON-Series switch-port profiles, see the *Dell EMC SmartFabric OS10 User Guide Release 10.5.0*.

NTP server configuration

This information explains how to set up network time protocol (NTP) to synchronize timekeeping among a set of distributed time servers and clients. The protocol coordinates time distribution in a large, diverse network. NTP clients synchronize with NTP servers that provide accurate time measurement. NTP clients choose from several NTP servers to determine which offers the best available source of time and the most reliable transmission of information.

NOTE: The NTP server configured on SFD should be on premise (located in the same data center as SFD), and reachable by SFD. Using NTP servers (such as time.google.com) that are not on premise or need Internet access for SFD to interface with is not recommended.

To get the correct time, OS10 synchronizes with a time-serving host. For the current time, you can set the system to poll specific NTP time-serving hosts. From those time-serving hosts, the system chooses one NTP host to synchronize with and acts as a client to the NTP host. After the host-client relationship establishes, the networking device propagates the time information throughout its local network.

For complete information about NTP, see the *Dell EMC SmartFabric OS10 User Guide*.

Enable NTP

NTP is disabled by default. To enable NTP, configure an NTP server where the system synchronizes. To configure multiple servers, enter the command multiple times. Multiple servers may impact CPU resources.

Enter the IP address of the NTP server where the system synchronizes in CONFIGURATION mode.

```
OS10(config)# ntp server ip-address
```

View system clock state

```
OS10(config)# do show ntp status
system peer:          0.0.0.0
system peer mode:     unspec
leap indicator:       11
stratum:              16
precision:            -22
root distance:        0.00000 s
root dispersion:      1.28647 s
reference ID:         [73.78.73.84]
reference time:       00000000.00000000  Mon, Jan  1 1900  0:00:00.000
system flags:         monitor ntp kernel stats
jitter:               0.000000 s
stability:            0.000 ppm
broadcastdelay:       0.000000 s
authdelay:            0.000000 s
```

View calculated NTP synchronization variables

```
OS10(config)# do show ntp associations
     remote           local    st poll reach  delay   offset    disp
==============================================================================
 10.16.150.185   10.16.151.123 16 1024    0 0.00000  0.000000 3.99217

OS10# show ntp associations
     remote           local    st poll reach  delay   offset    disp
==============================================================================
 10.16.150.185   10.16.151.123 16 1024    0 0.00000  0.000000 3.99217
```

6

First-time setup

This information explains what you must do if setting up SmartFabric Director for the first time.

Topics:

Installation using vCenter 6.7

This information describes how to import the SmartFabric Director OVA file into the content library, then create a virtual machine (VM). It is recommended that SFD is installed on a server which is part of your infrastructure rack, and is different from workload servers.

Download and install OVA

You can add items to a content library by importing files from your local system. You can import an OVA package to use as a template for deploying virtual machines.

1. Download the OVA from DDL or the VMware Solution Exchange, then store the OVA image locally or on a server.
2. Select **Hosts and Domains**, select the domain that the plug-in must manage, then select **Action** > **Deploy OVF Template**.
3. Select **Local file**, click **Choose Files** and select `OMNI.ova` from a local source, then click **Next**. You can use either a URL or a local file.

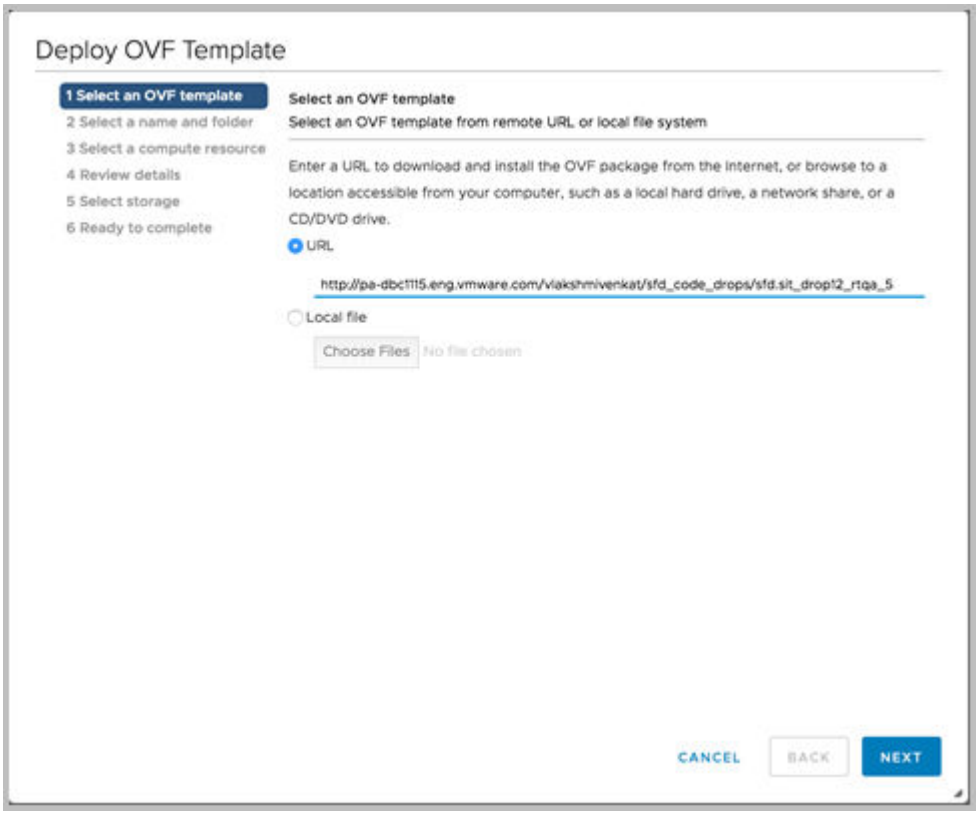

4. Select a name and folder for the VM, then click **Next**.

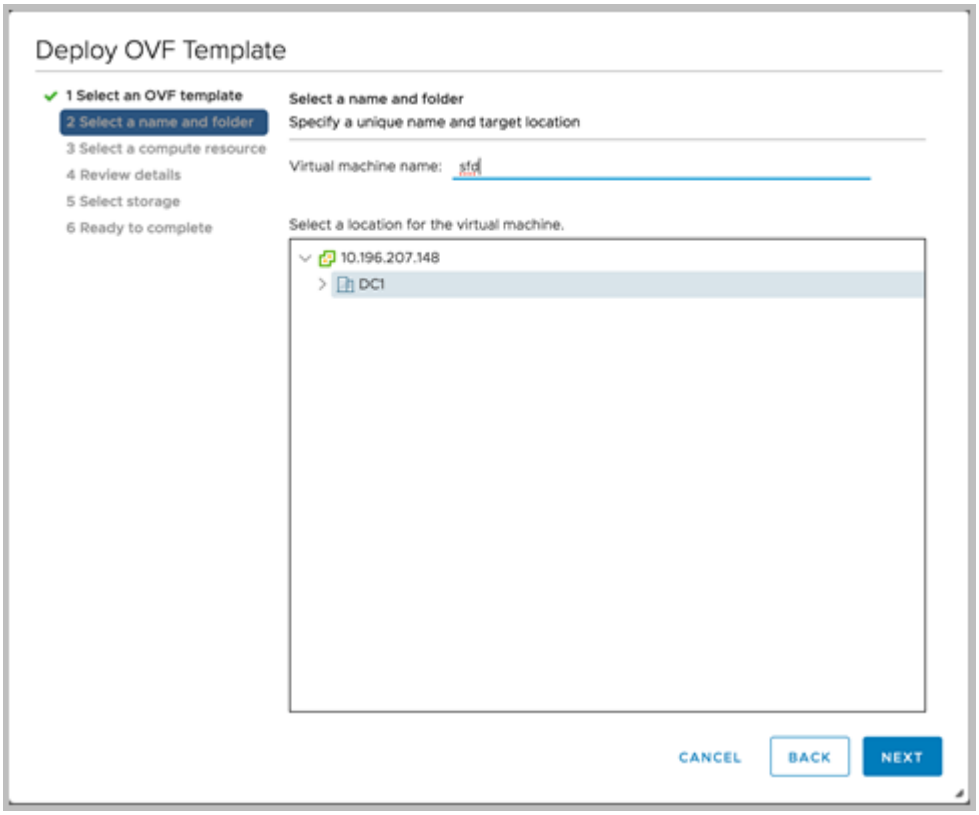

5. Select the destination compute resource, then click **Next**.

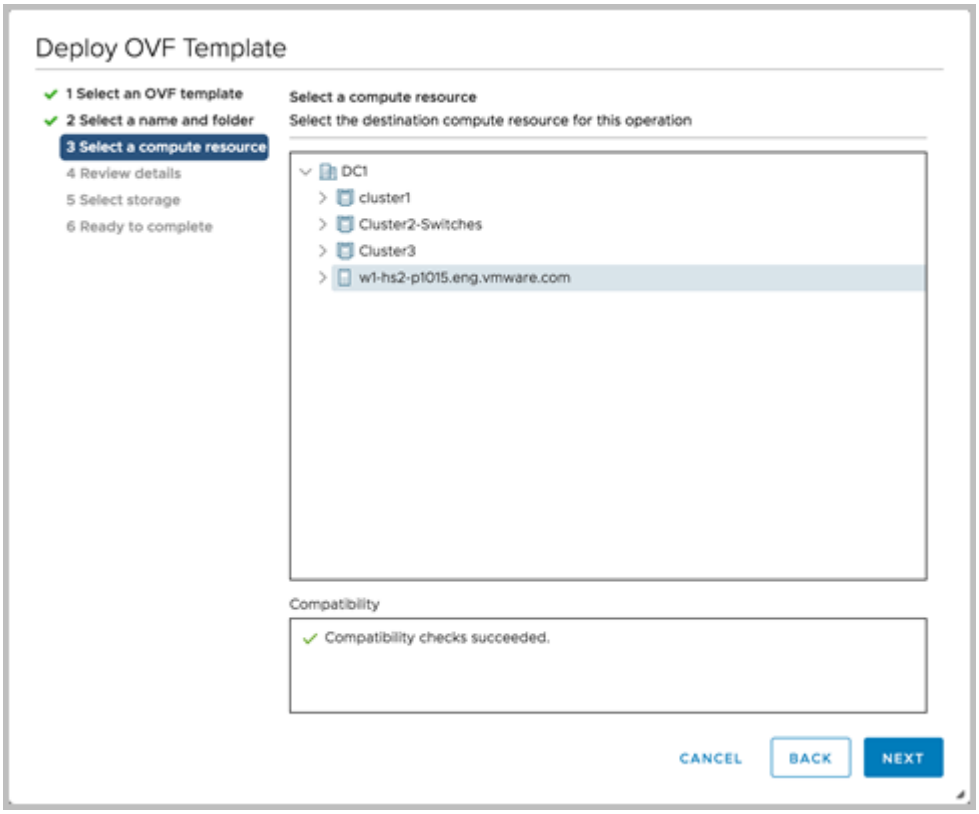

6. Review and verify the template details, then click **Next**.

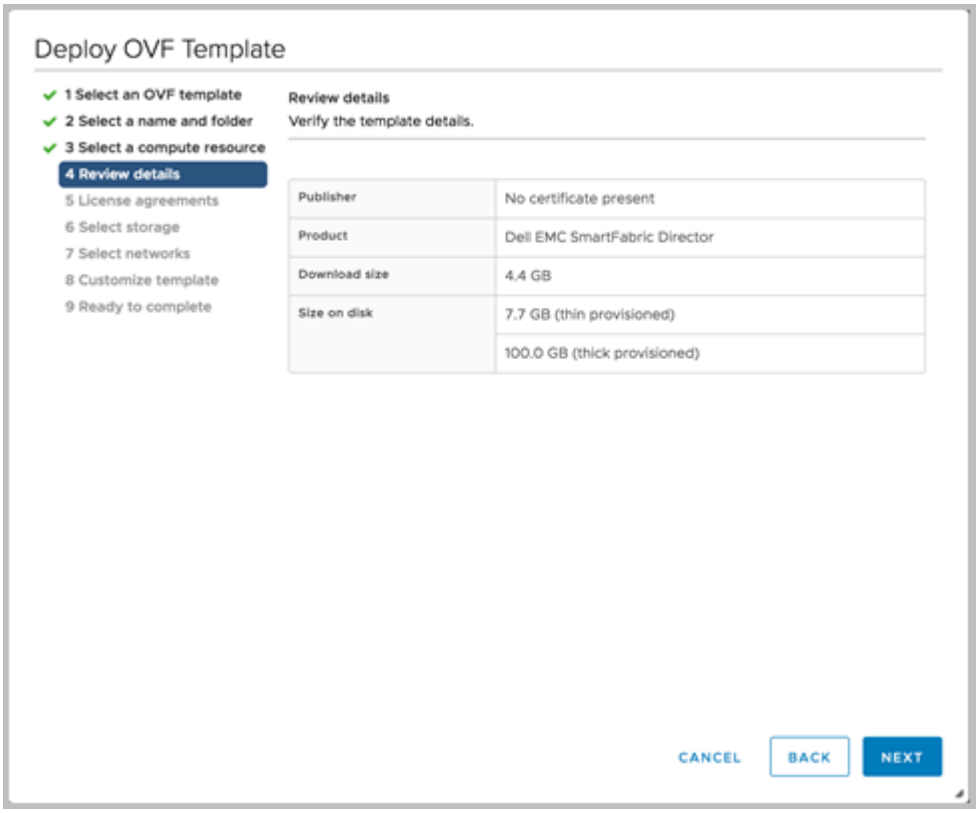

7. Accept the user license agreement (EULA), then click **Next**.

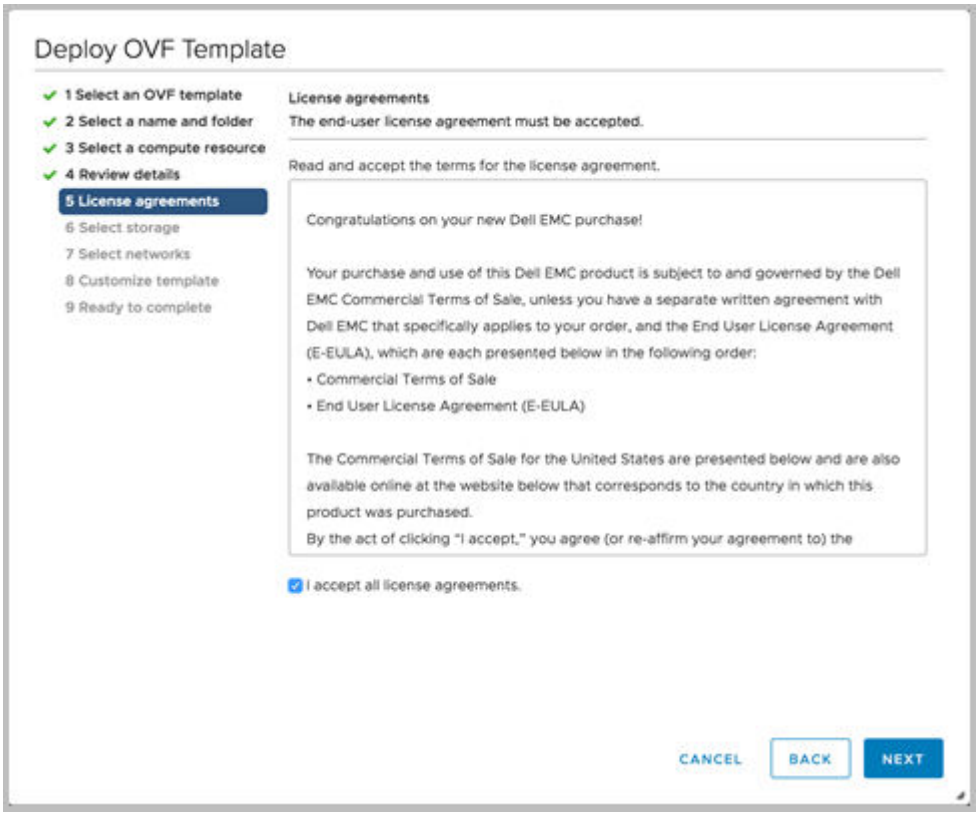

8. Select the data store to store the configuration and datafile, then click **Next**.

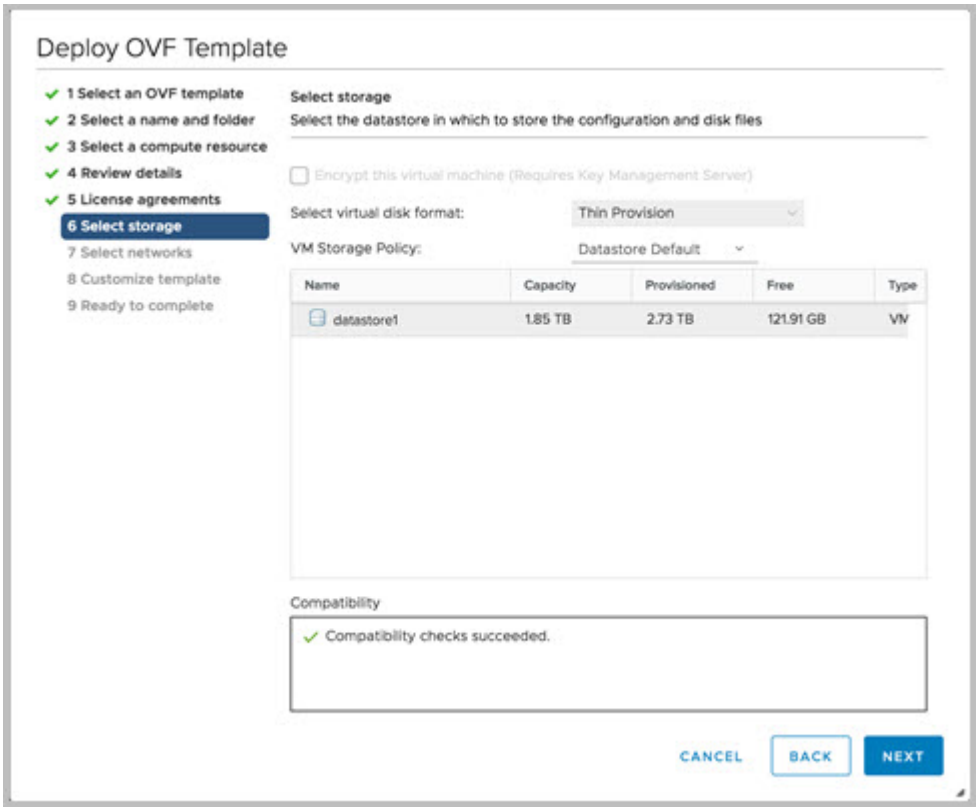

9. Select a destination network for each network source, then click **Next**. The default VLAN ID for this network is 3939. The vCenter Server network must be connected to the port group where the vCenter is reachable for plug-in deployment of the VM.

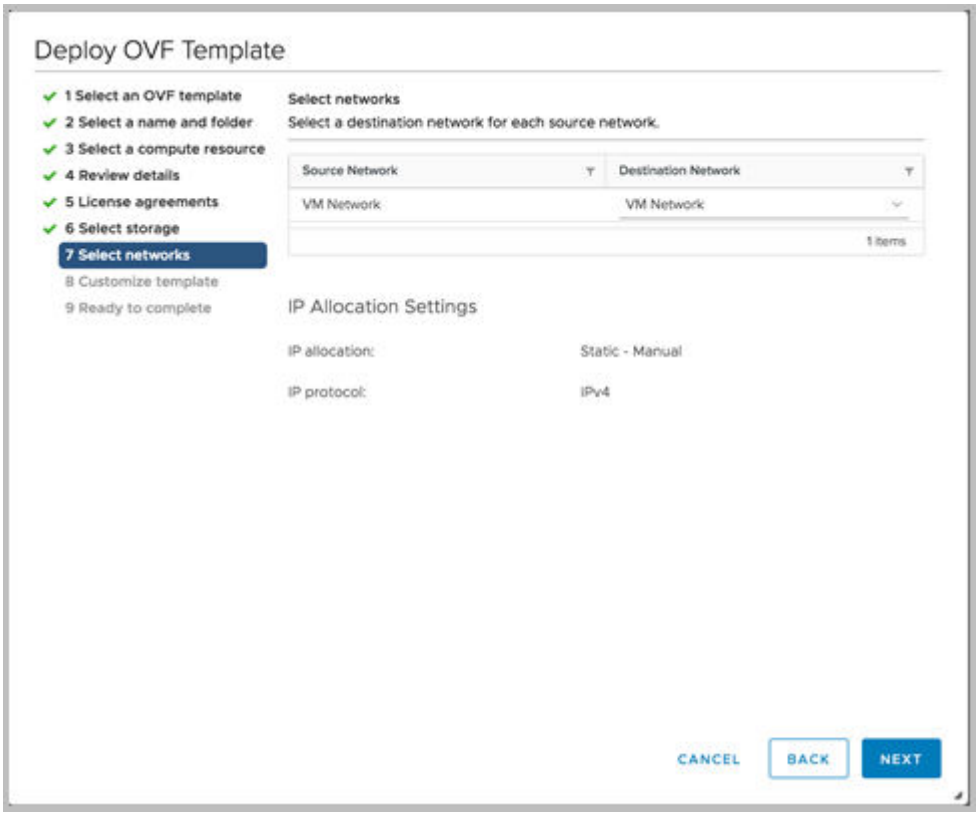

10. Enter the system name for the appliance, select the checkbox if SSH access is enabled, list the NTP servers (space separated), enter the Domain Name Server, then click **Next**. The Service Tag can be added later as part of the system settings after logging into SFD.
11. Select **Networking Properties** to customize the template.

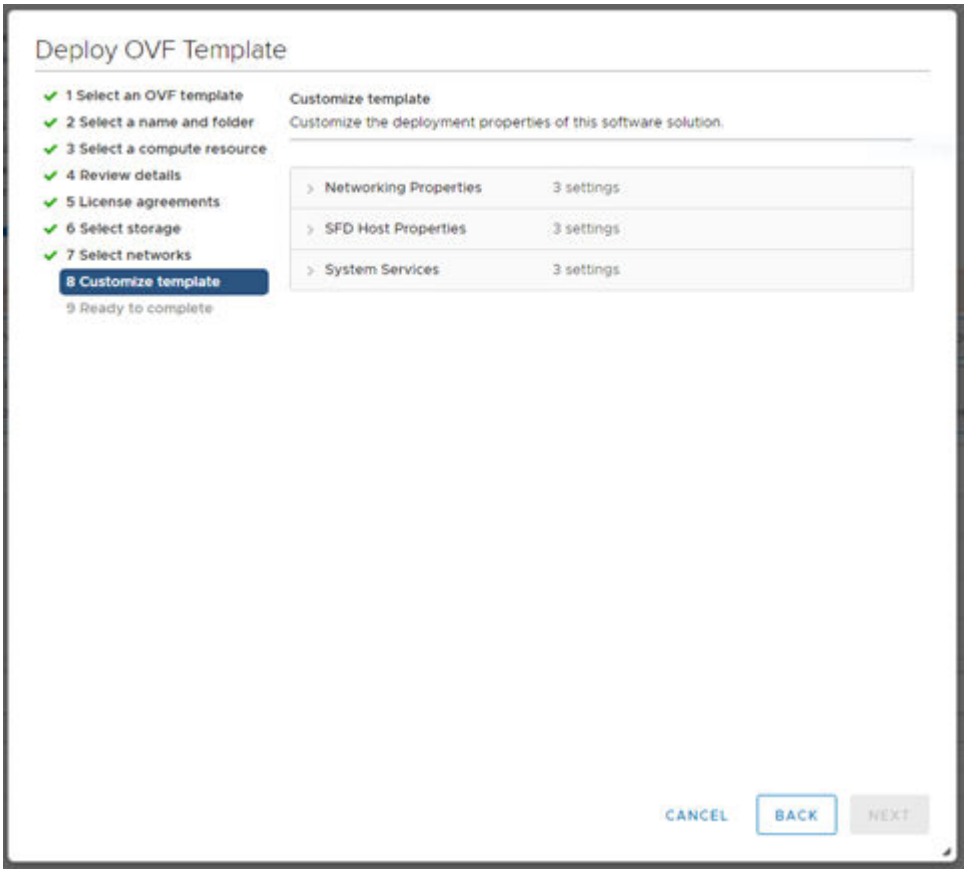

12. Enter the IPv4 address for this interface, the netmask, and the default IPv4 gateway address for this VM.

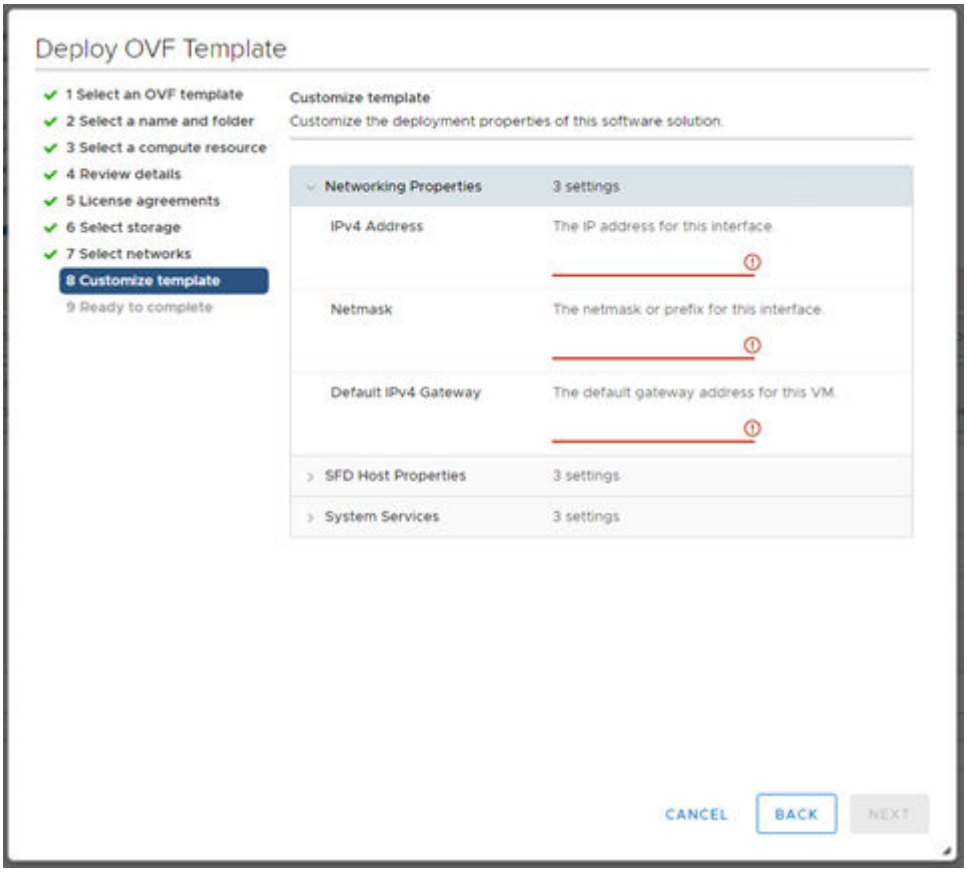

13. Select **SFD Host Properties**, then enter the password for the SFD host.

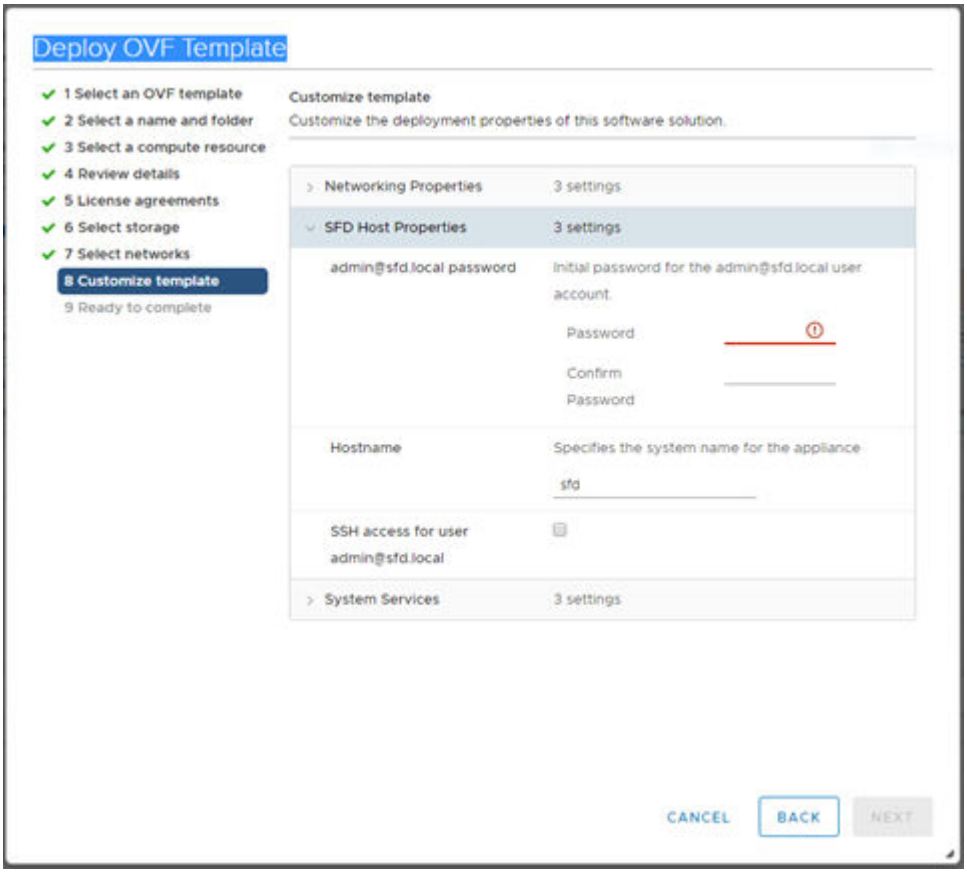

14. Select **System Services**, enter the list of NTP servers separated by a space, then click **Next**.

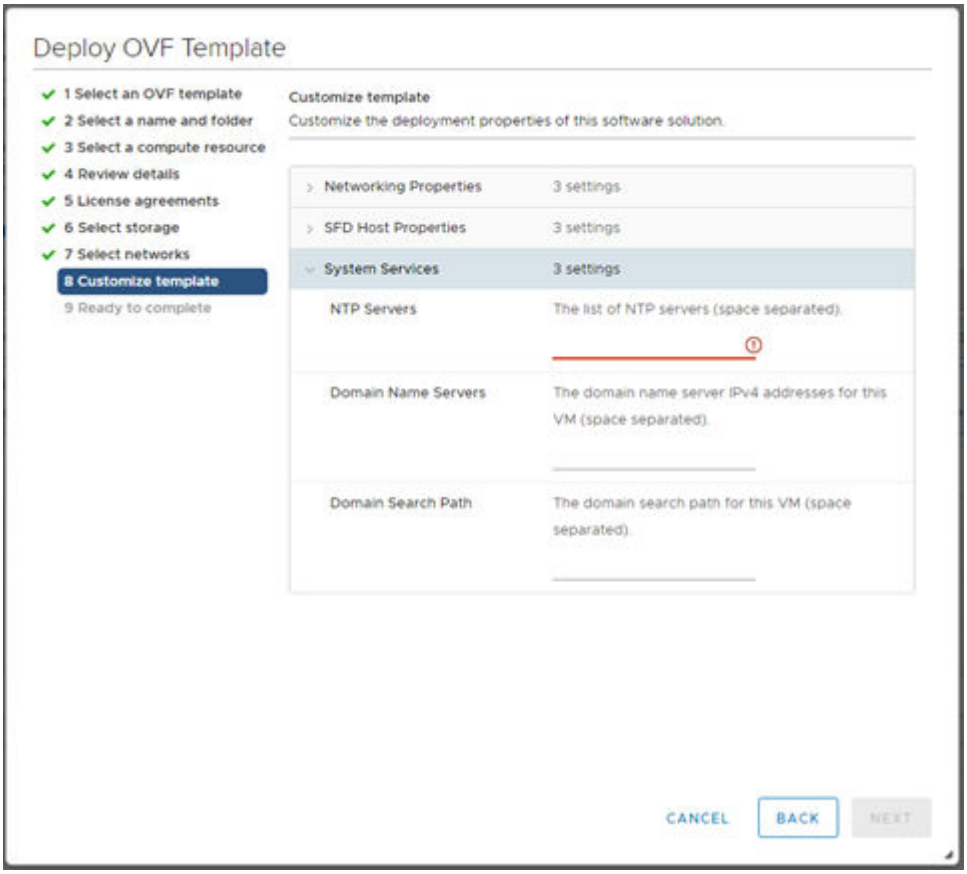

15. Click **Finish** to start creation of the VM, then power on the VM.

NOTE: Once installation finishes, it may take 7 to 12 minutes for the SFD VM to be fully operational.

Installation using vCenter 6.5

This information describes how to modify the .vmx file for a successful SmartFabric Director installation on vSphere 6.5. You must edit the .vmx file and comment out the nvram location setting.

NOTE: If you are using vSphere 6.7, go to the next section. If you are using vSphere 6.5, you must modify the nvram setting before creating a VM.

1. Shut down the SFD virtual machine.
2. Download the .vmx configuration file from the VM folder, then open the file in a text editor.
3. Comment out the nvram setting.

```
.encoding = "UTF-8"
config.version = "8"
virtualHW.version = "10"
pciBridge0.present = "TRUE"
svga.present = "TRUE"
pciBridge4.present = "TRUE"
pciBridge4.virtualDev = "pcieRootPort"
pciBridge4.functions = "8"
pciBridge5.present = "TRUE"
pciBridge5.virtualDev = "pcieRootPort"
<snip>
#nvram = "ovf:/file/file2" <- The 'nvram' entry must be commented out
```

4. Save the changes, replace the file on the VM folder, then exit the text editor.
5. Start the VM.

NOTE: If SmartFabric Director is restarted and if any changes have been made to the switches (such as switches reloaded) during this time, it is recommended to reload the switch again after SFD is up. This ensures that the complete configuration is redownloaded to the switches.

Log in to SmartFabric Director

This information explains how to log in to SmartFabric Director.

1. Open a browser window, then enter the IP address that is specified during installation of the SFD VM in vCenter.
2. Enter **admin@sfd.local** for the username and the password that is configured during VM deployment, then click **Login**.

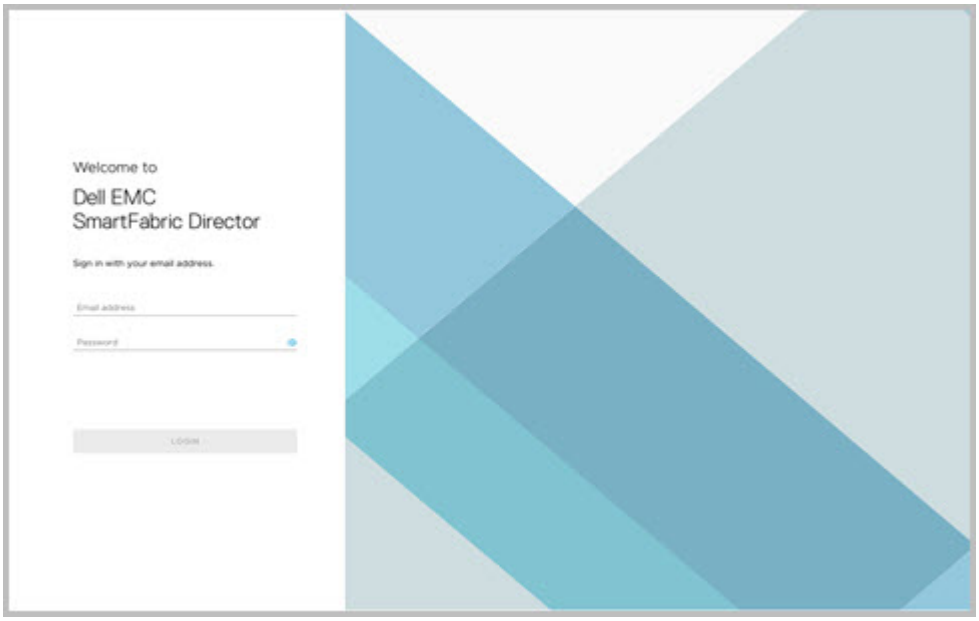

If the system administrator has enabled Active Directory, enter your AD username and password, then click **Login**.

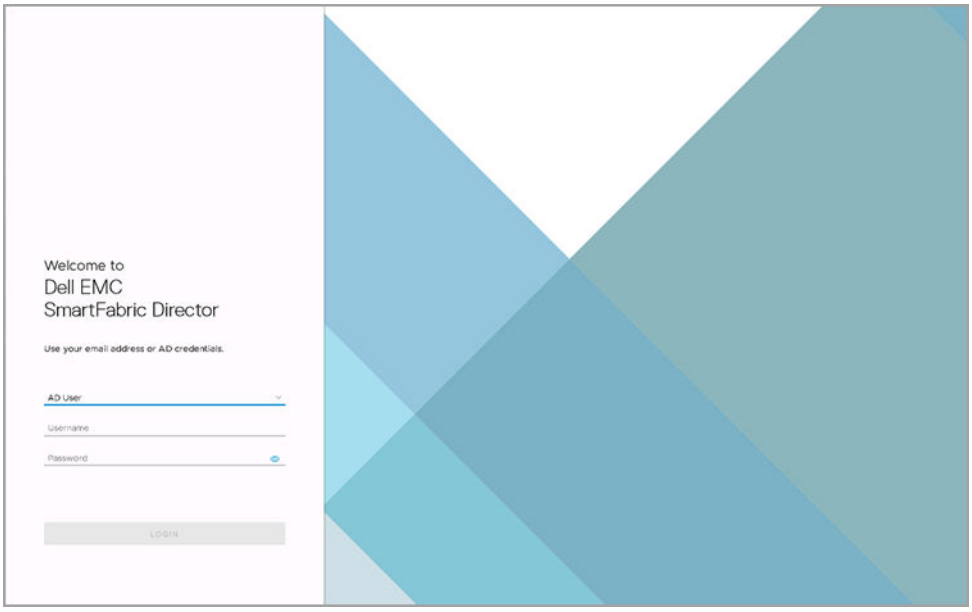

Specify system settings

This information describes the system settings for SmartFabric Director.

About

The About tab displays the SmartFabric Director software version and the Software Service Tag. This information is helpful when upgrading the software.

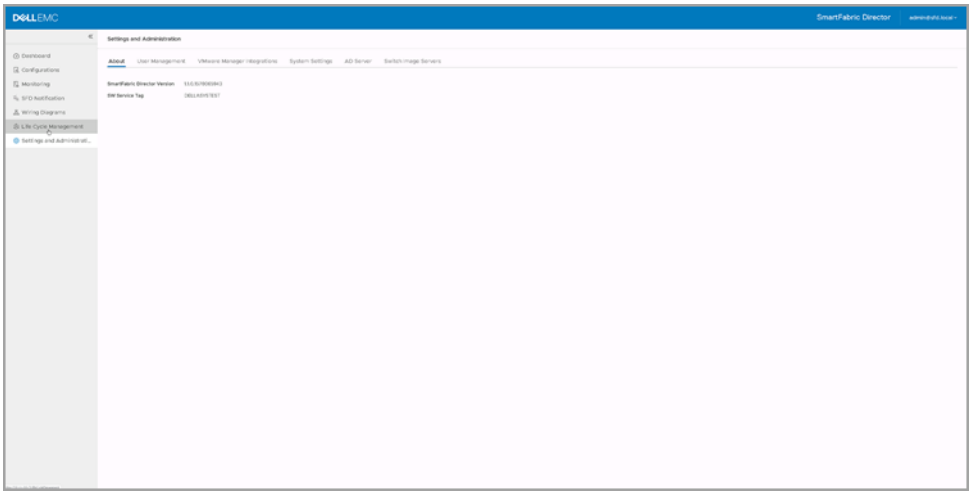

VMware Manager integrations

NOTE: You can only have one active VMware Manager integration with SmartFabric Director.

1. Select **Settings and Administration > VMware Manager Integrations**.
2. Click **Add VMware Manager** to configure the VMware Manager integration.
3. Enter the required connection information for the VMware Manager, then click **OK**.

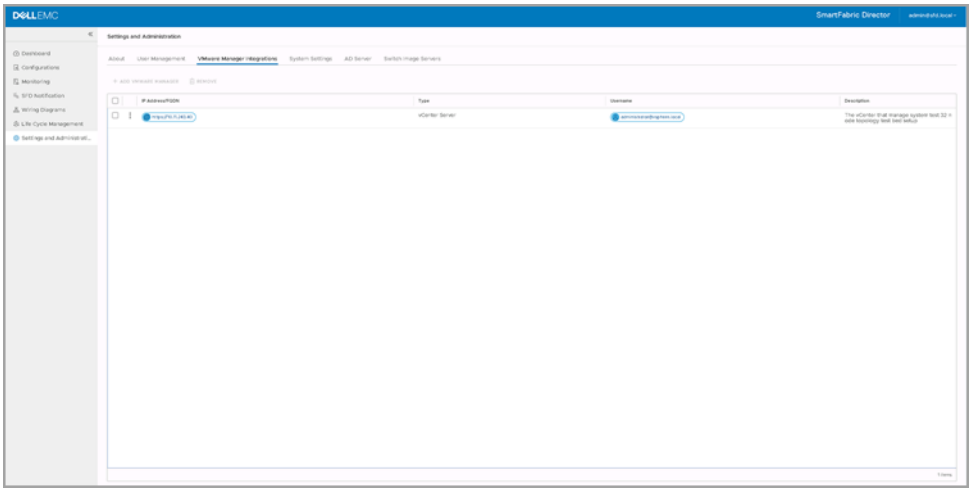

4. (Optional) To remove the current VMware Manager connection, click **Remove** and follow the steps.

   NOTE: If you remove the current VMware Manager Integration, you must add a new VMware Manager connection to OMNI.

System settings

1. Select the **Settings and Administration** icon from the left, then select **System Settings**.

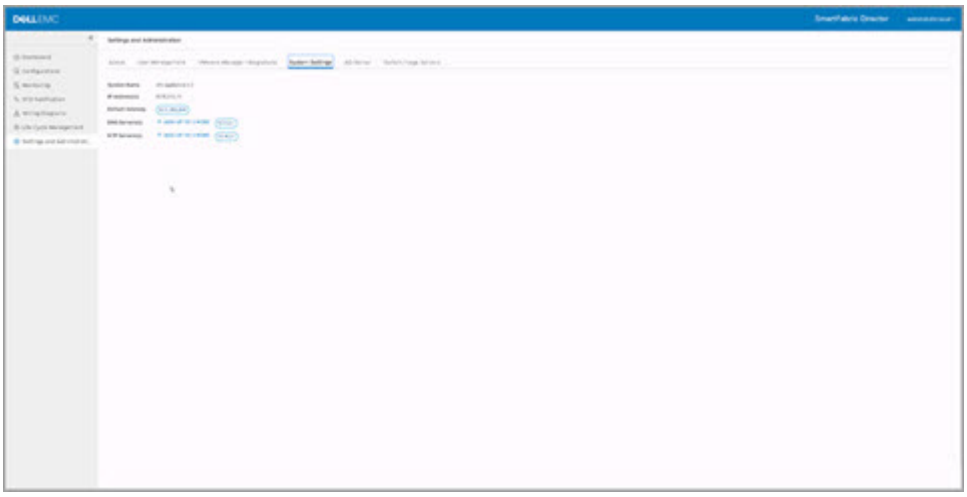

2. Click **Add DNS Server**, enter the **IPv4 address**, then click **Save**.

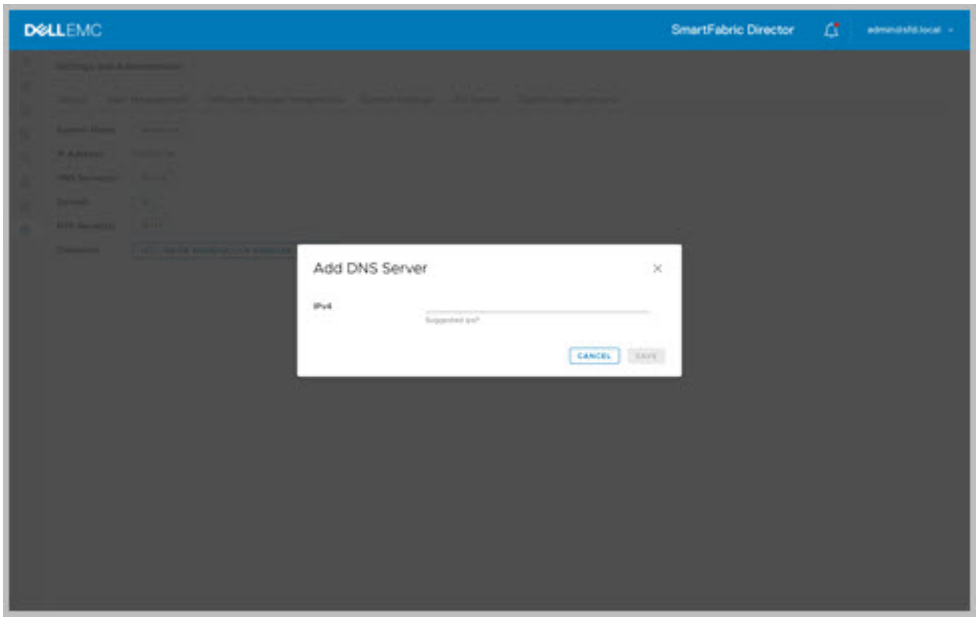

AD server

1. (Optional) Select the **AD Server** tab to specify an Active Directory Server.

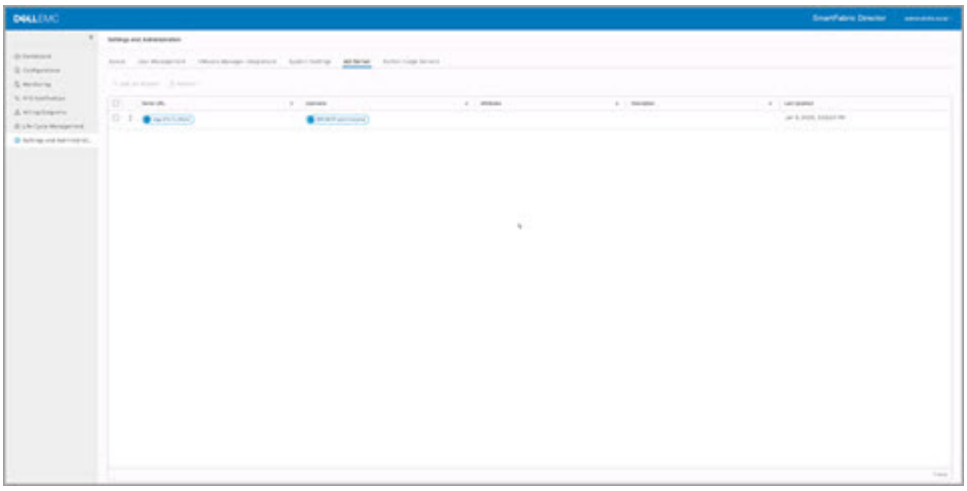

2. (Optional) Click **Add AD Server**, enter the server URL, username (admin), password, optional attributes, and optional description, then click **Add**.

   NOTE: If you do not set up an Active Directory Server, you must create user accounts.

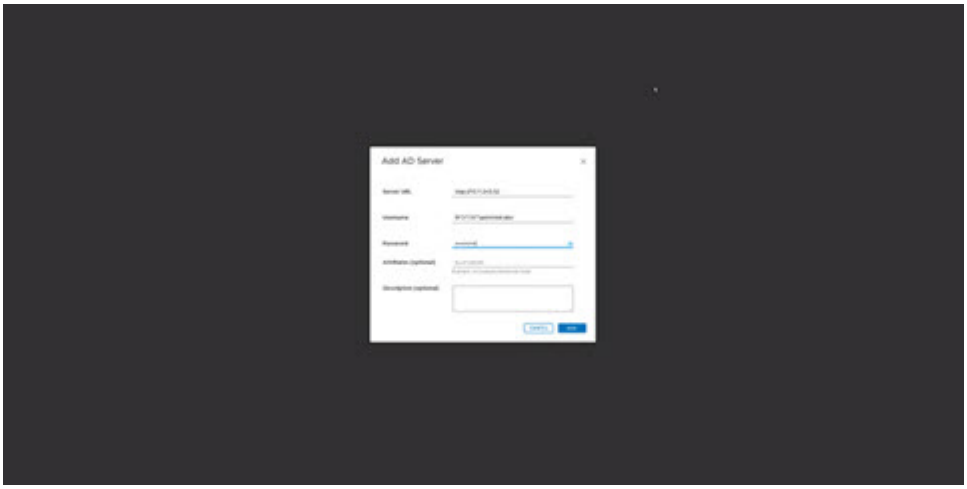

You can also edit or delete a previously added AD server.

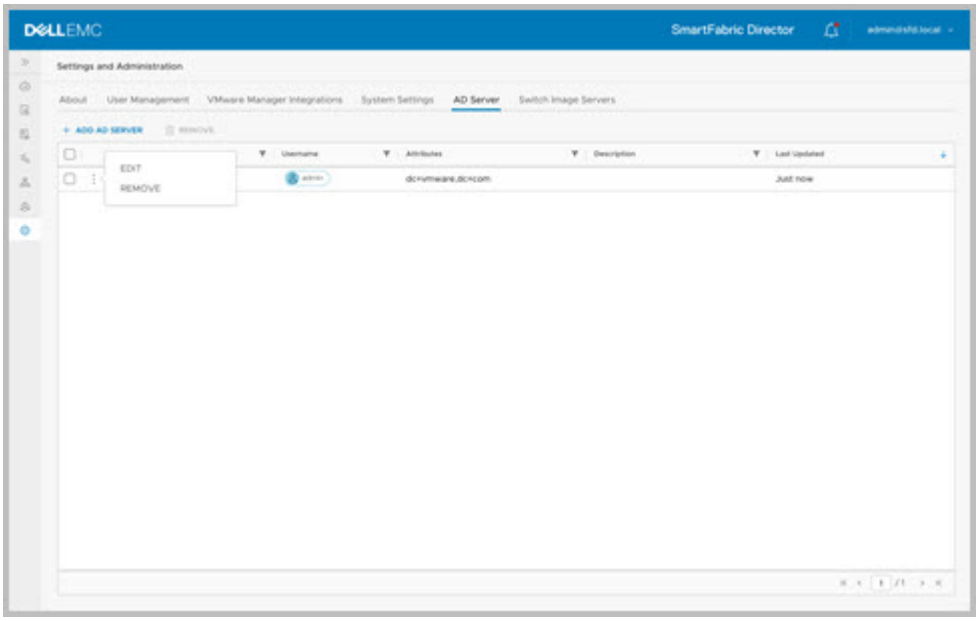

Switch image servers

1. Select **Switch Image Servers**, then click **Add image server**.

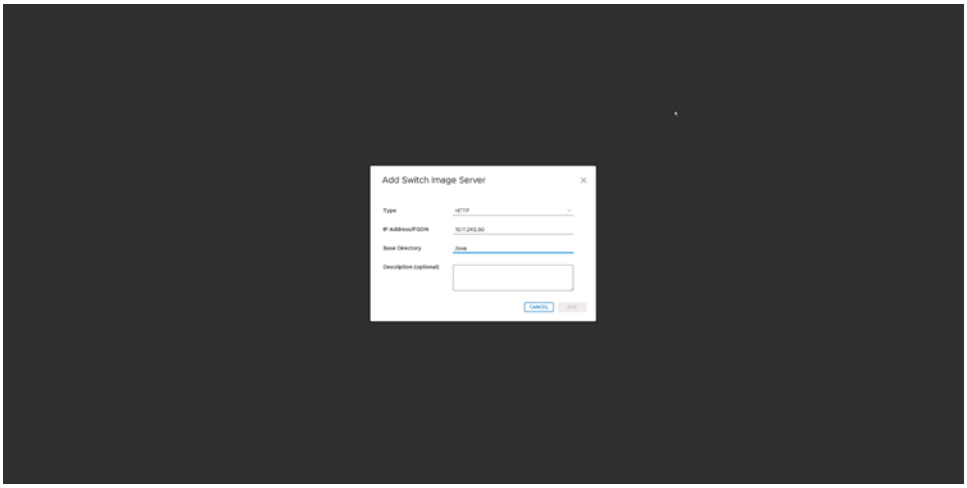

2. Select the image transport type, enter the IP address/FQDN, enter an option description, then click **Add**. The new image server information displays.

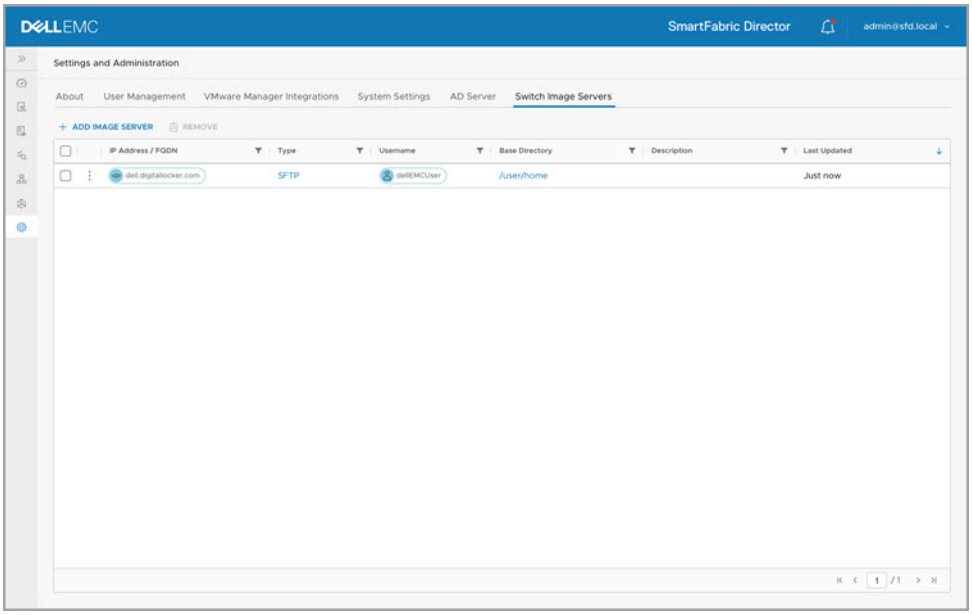

3. (Optional) Select the image server checkbox, then click **Remove** to delete the image server.

Upload service tag to SFD

1. Sign into DDL using your account credentials (see Download SFD image).
2. Download the SFD license file (XML) to your local file system.
3. Select Settings and Administration, then click **Upload**.

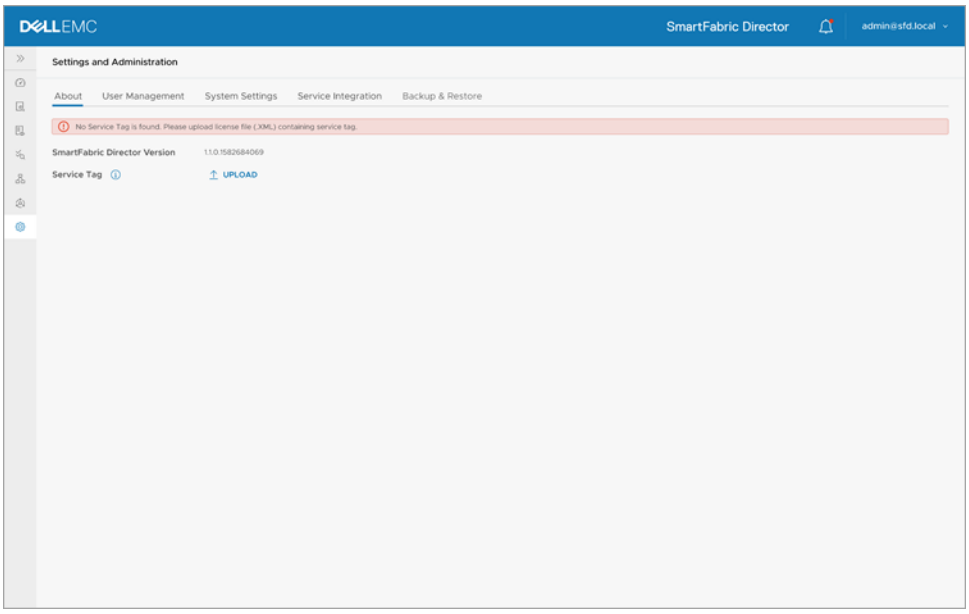

4. Click **Browse** to locate the XML license file, then click **Upload**.

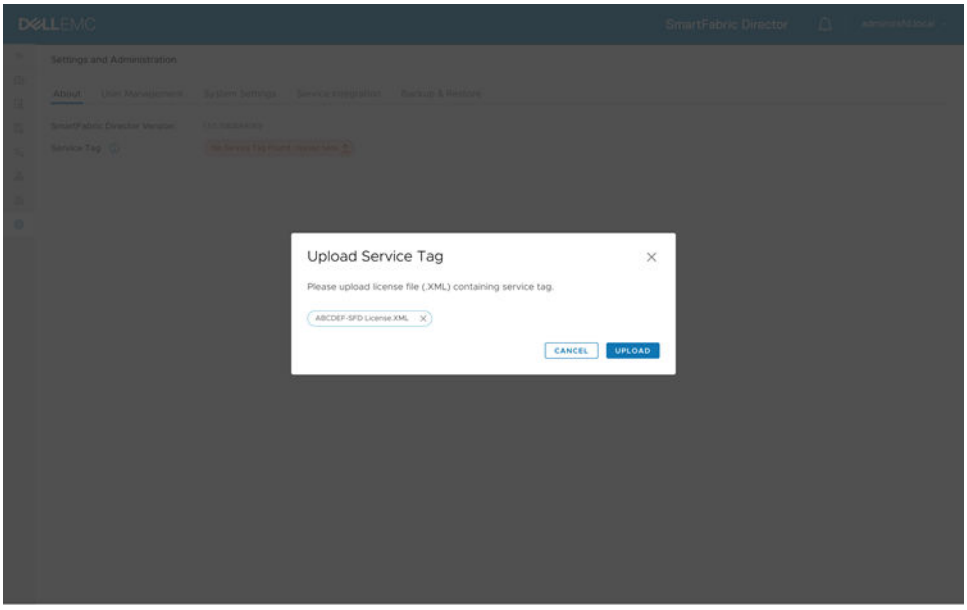

The upload service tag progress displays.

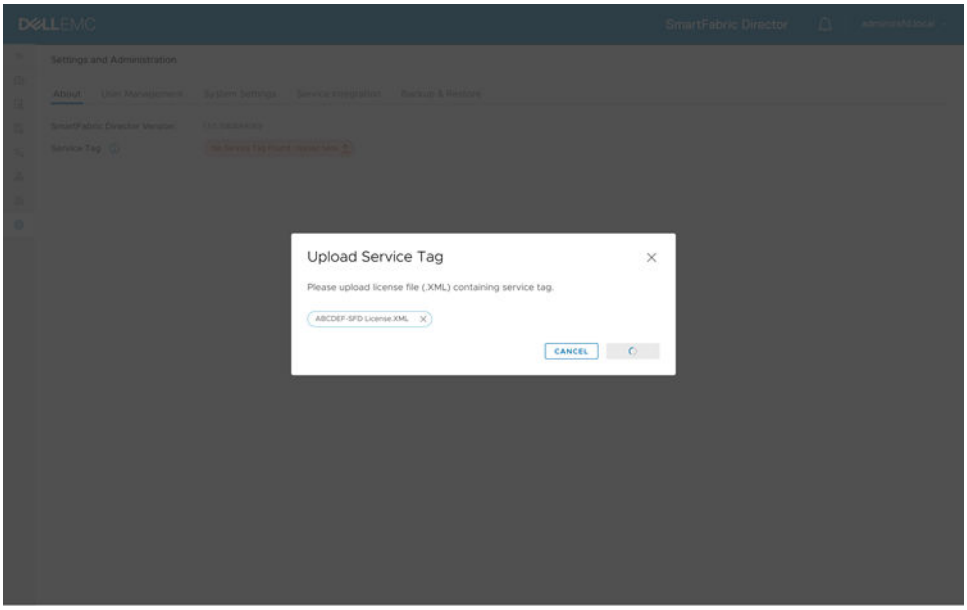

When the upload completes, the progress updates. Click the checkbox to finish.

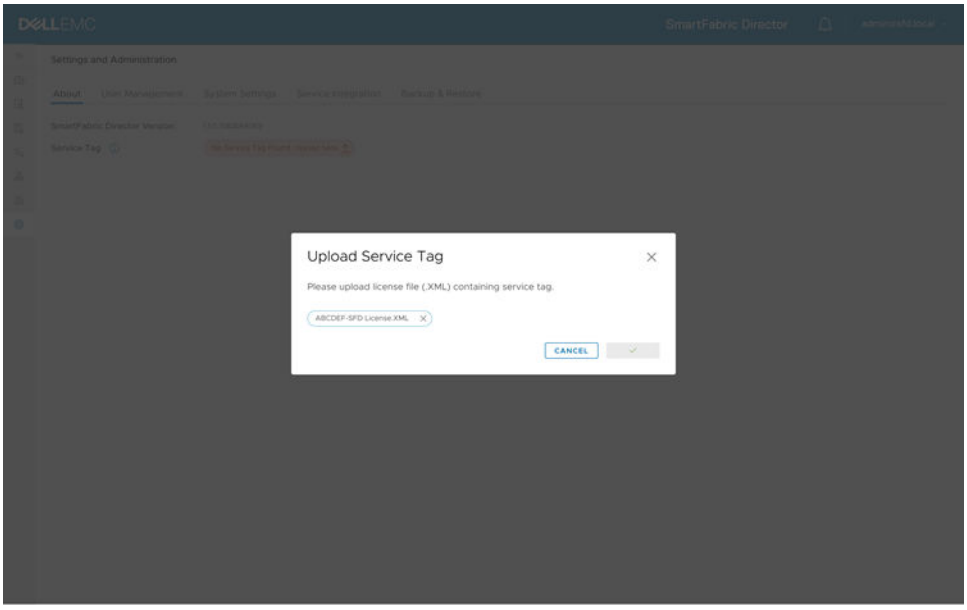

SFD extracts the service tag from the license file, then updates the About page.

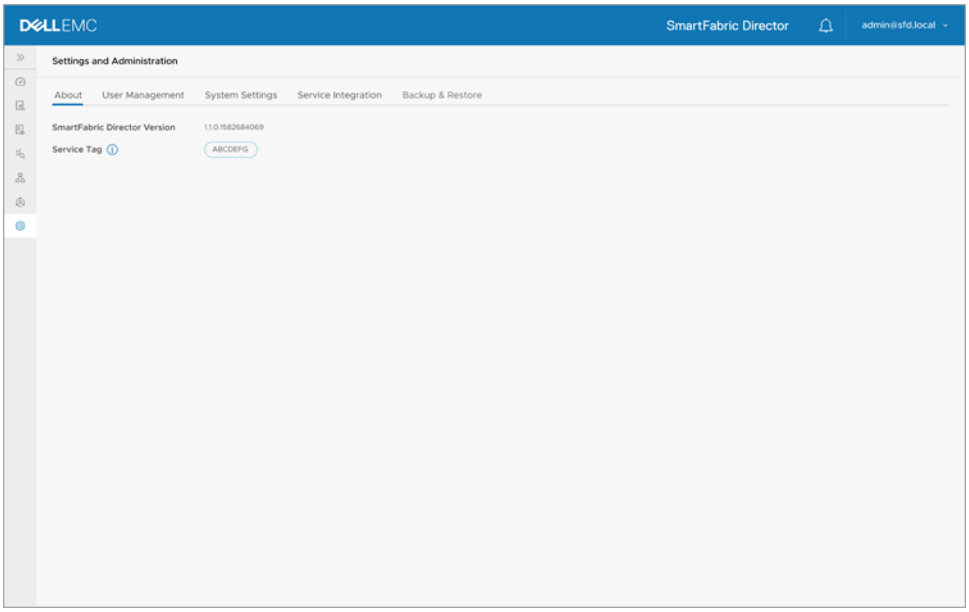

Create user accounts

This information describes how to add user accounts to SmartFabric Director. You can add local users through SmartFabric Director, or you can use the Active Directory Server. You can enable or disable SSH access for a local user, activate or deactivate, and remove a local user.

NOTE: If you have setup an Active Directory Server, you do not need to add local users.

1. Select the **User Management** tab to add local users.

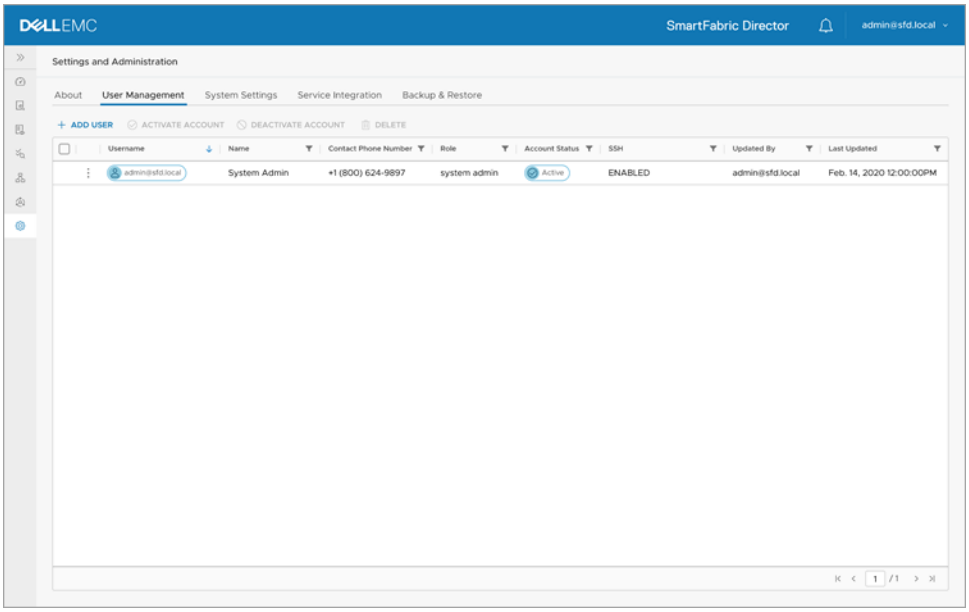

2. Click **Add local user**, then enter the new user email address, password twice, first and last name, optional phone number, select the user role, then click **Add**.

   NOTE: Each new local user has admin role privileges automatically, and SSH status is enabled by default.

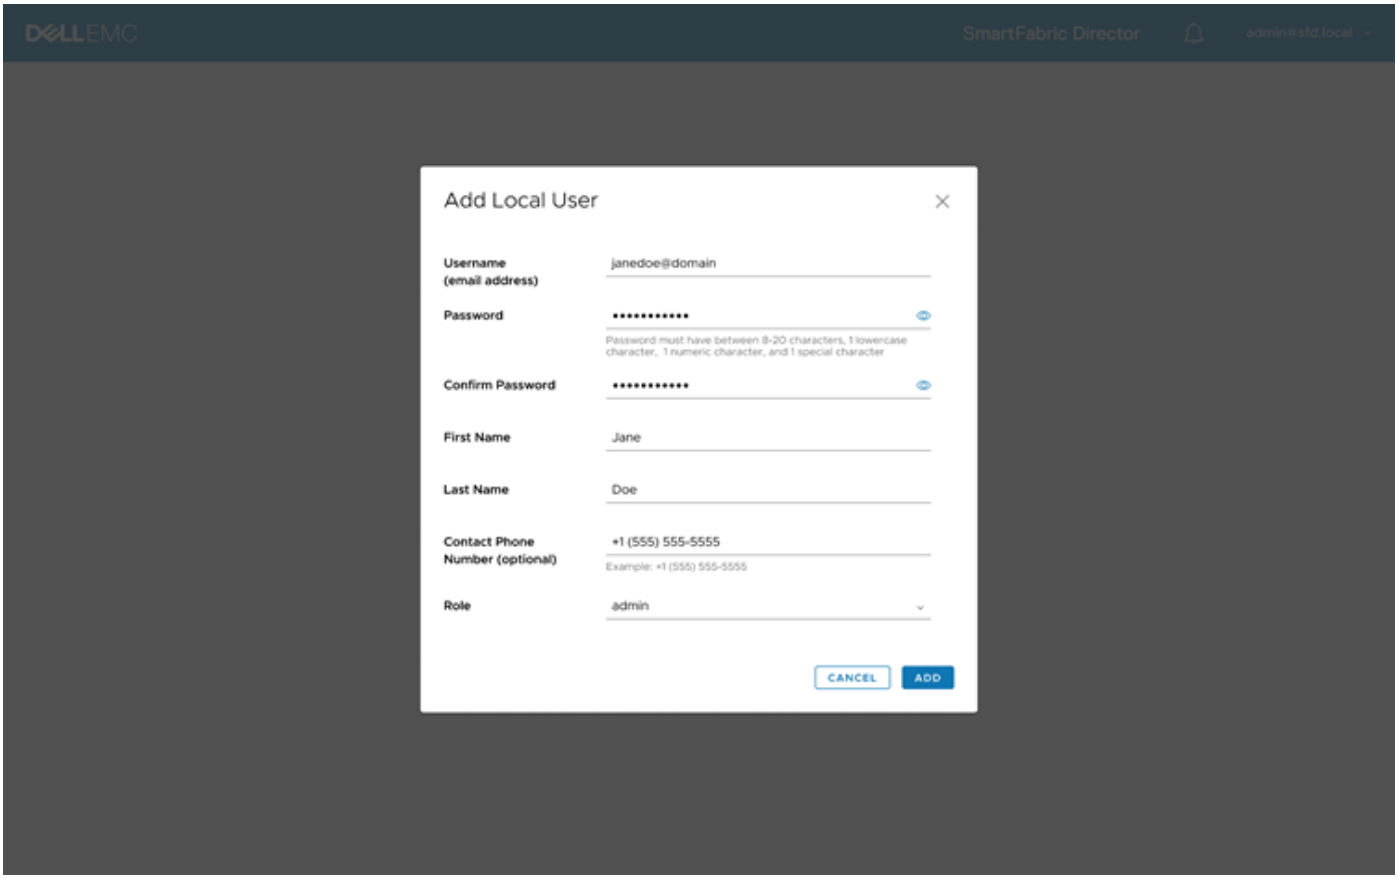

3. (Optional) Continue adding new users, or you can select **Edit** to modify any user profile, or select **Remove** to delete the user account.

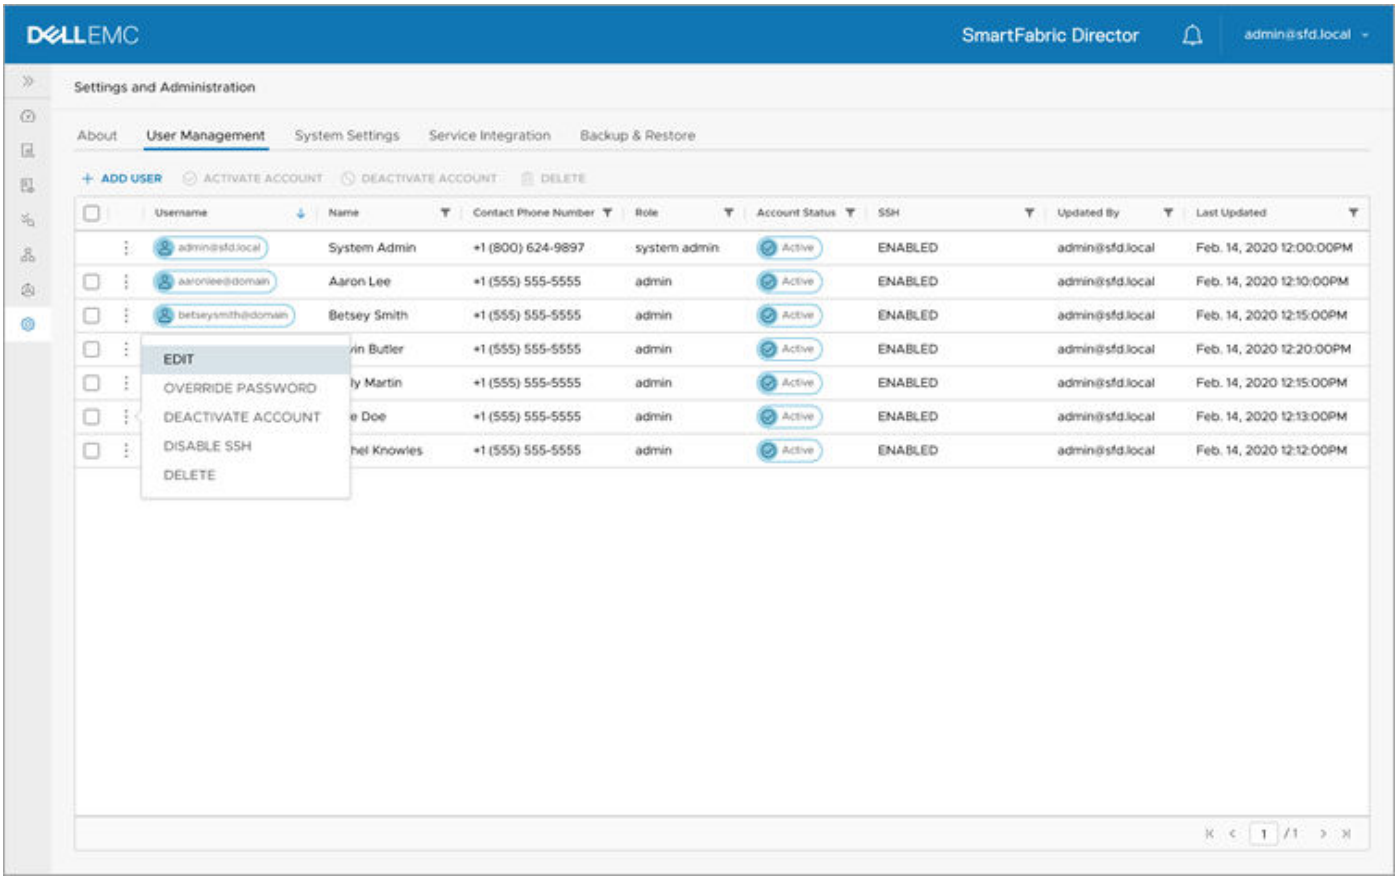

4. (Optional) Modify the user profile, then click **Save**.

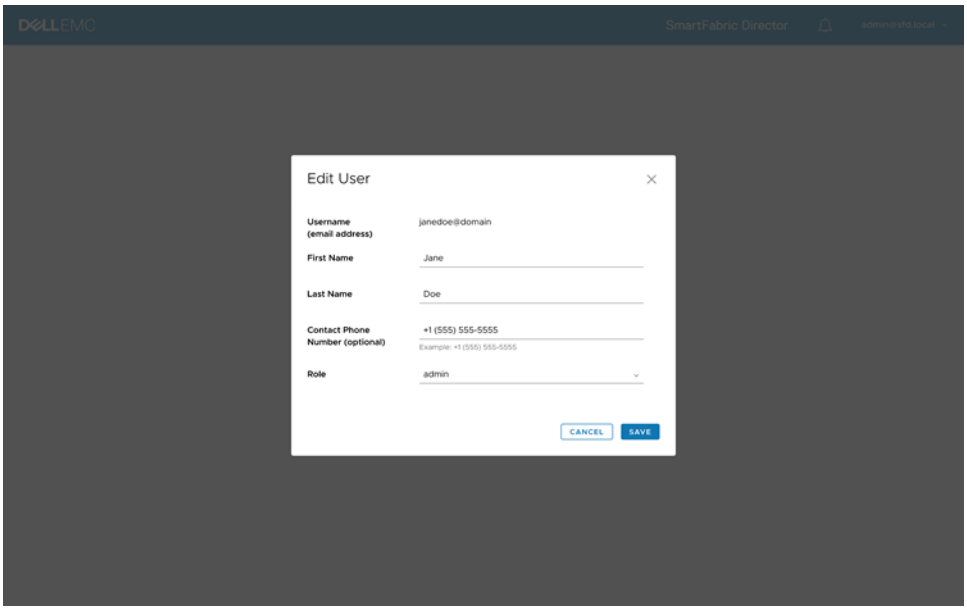

7

Using SmartFabric Director

This information describes how to build, define, and deploy a data center SmartFabric. After completion of the SFD bootstrap and integration into your data center network operations, you are now ready to build, define, and deploy a SmartFabric. The steps outline a deployment where the operator specifies the SmartFabric for the first time, starting with a clean state.

Topics:

Import fabric wiring diagram

This information describes how to use the fabric definition screens to import a fabric wiring diagram. Fabric definition describes the switches, their roles (spine, leaf, or edge leaf), and the wiring diagram of how these switches interconnect.

You must specify the reachability information, such as the Management IP and credentials (username and password) of each switch, so that SFD can connect to the switches.

NOTE: The fabric wiring diagram must be edited manually using a text or JSON editor. Using Fabric Design Center enables automatic generation of the JSON file.

It is assumed that the switches have been racked, stacked, and connected as per the wiring diagram. All switches must have the minimum version of SmartFabric OS10 10.5.0.4 installed, along with the base configuration to connect, and communicate through the gNMI and gNOI interfaces. For more information about the base configuration, see Configuration.

You must define the role of each switch, and the interface type such as interlink, host, or edge facing. It is expected that the switches are wired per the definition in the wiring diagram.

- Interlinks are switch ports that are used to connect a leaf switch to a spine switch
- Host interfaces are switch ports that are connected to host
- Edge interfaces (on an edge leaf) are switch ports that are connected to an external switch

Using Fabric Design Center

Dell EMC Fabric Design Center allows you to automatically create a JSON file based on your selections.

1. Open a browser, then go to fdc.emc.com and log in with your Dell Customer/Partner credentials.
2. Select the checkbox to agree to the terms of use, then click **OK**.
3. Select either *Designing network fabric for a customer opportunity* or *Trying Fabric Design Center*, then click **OK**.

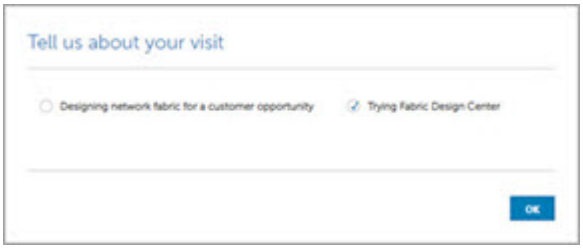

4. Select **Build-Your-Own-Network Design** at the top.
5. Create your Layer 2 or Layer 3 network, then click **Apply**. You can select the number of racks that are needed for the design. The default number of spines is calculated based on the default bandwidth to the rack.

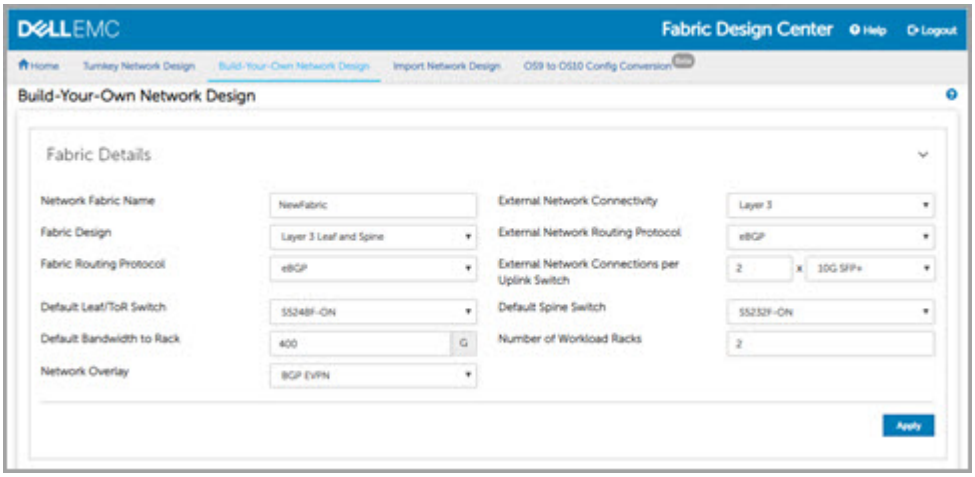

6. Verify the fabric design, then click **Next**. You can also click Edit and make any necessary changes to meet your requirements.

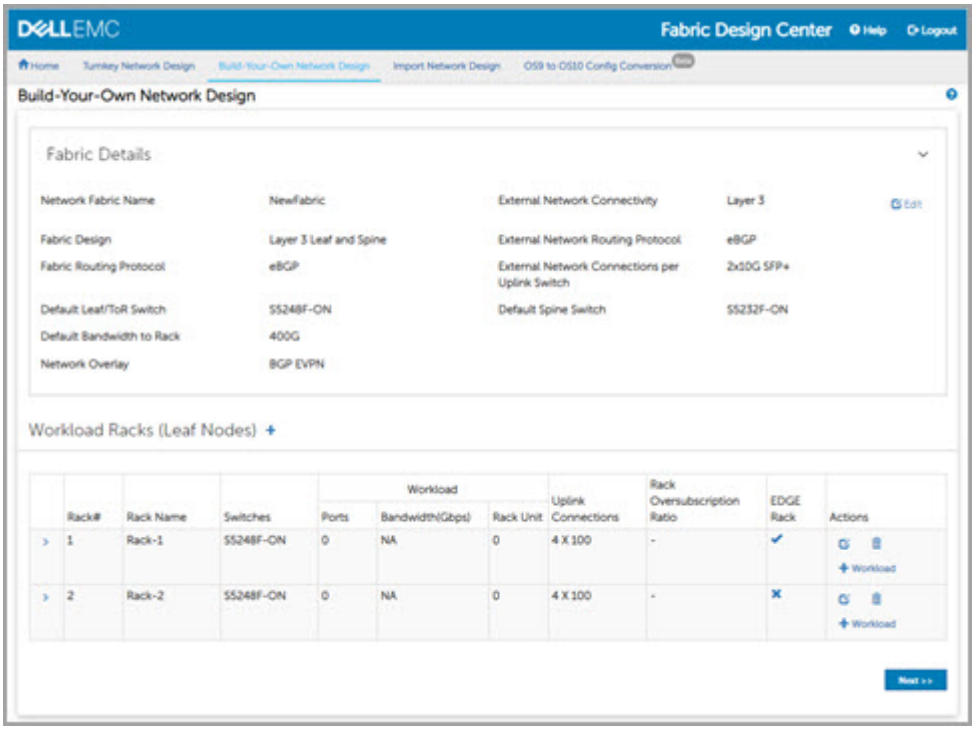

7. Select **SFD wiring** for the file format, then click **Download** to save the JSON file for importing into SFD.

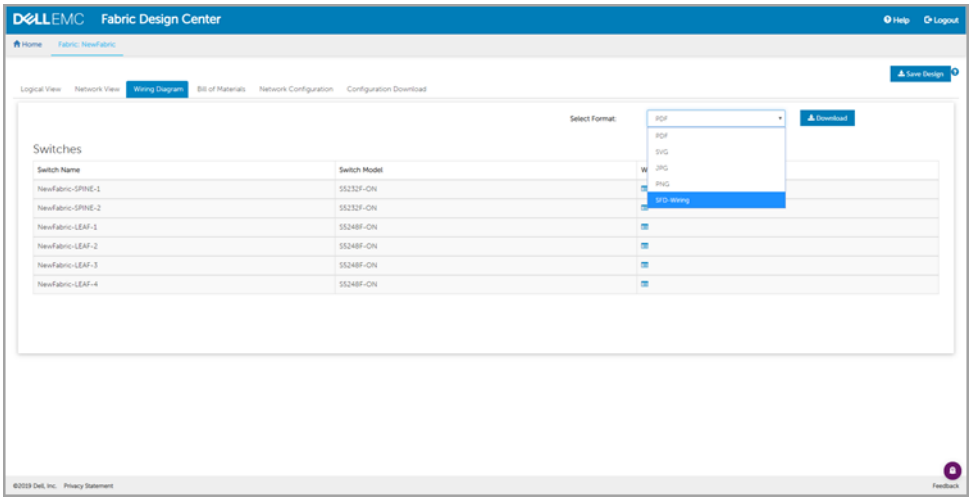

For complete information, see the *Dell EMC Fabric Design Center User Guide*.

Import a wiring diagram

You can import a JSON wiring diagram file through the user interface. On reading the JSON file, the user interface displays the fabric graph as described in the JSON file. You can download a JSON wiring template from DDL .

NOTE: Once the JSON file is imported or activated, any active intent of the fabric is no longer valid, and a new intent must be defined and submitted for approval.

1. Click **Import Fabric**.

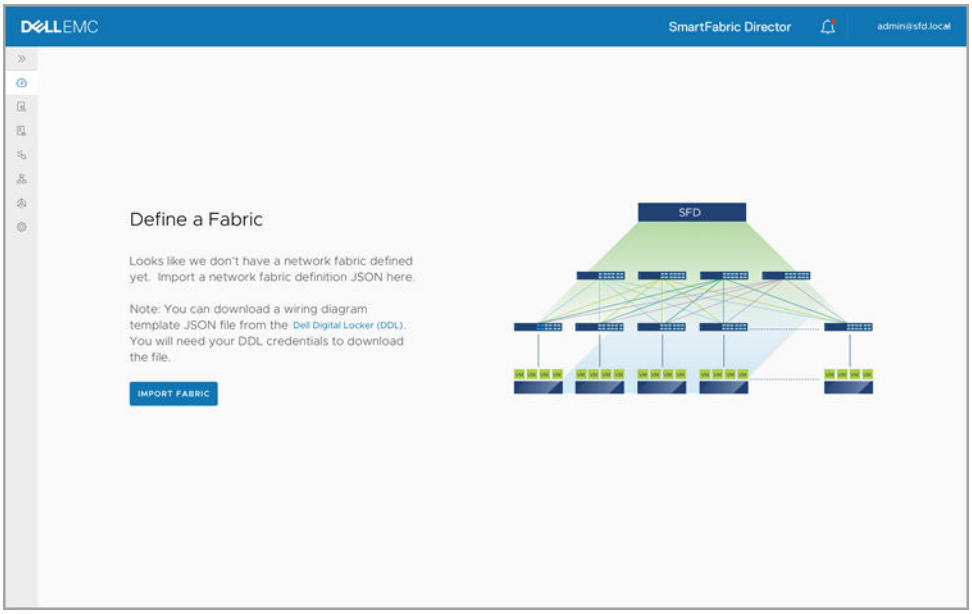

2. Go to the file location, then select the JSON file to import.

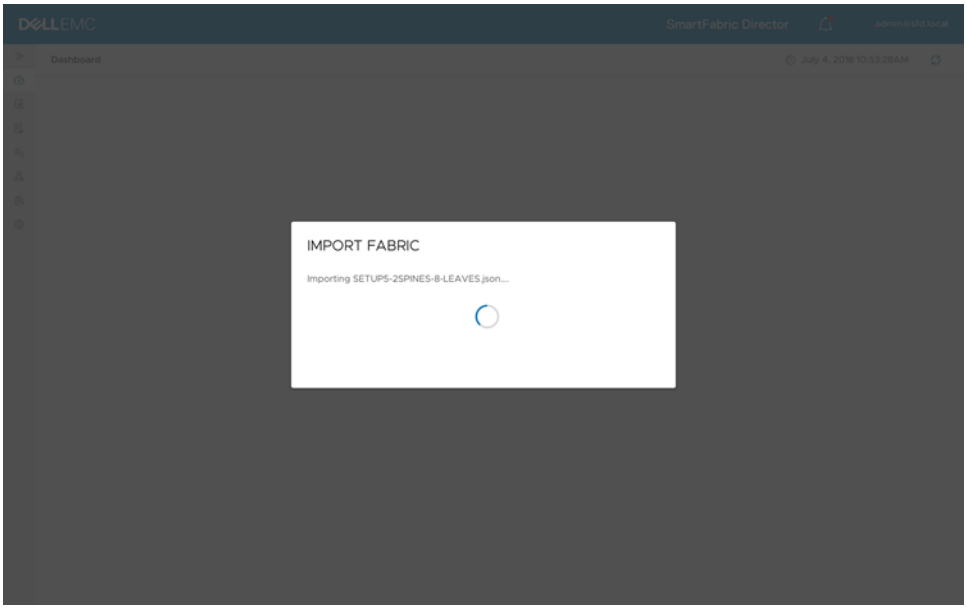

Wiring import success.

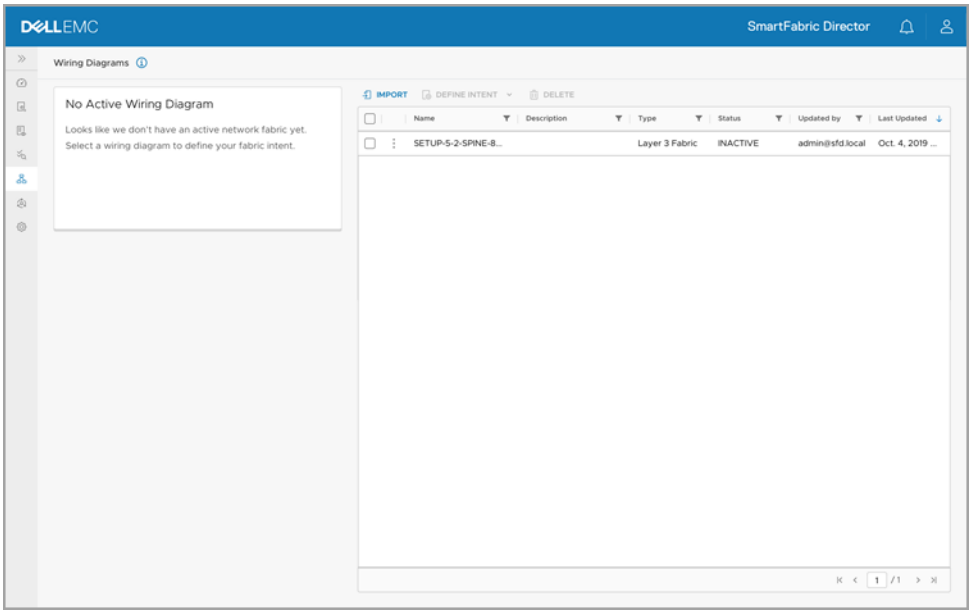

3. Select **Create New Intent** to define and associate a fabric intent with the wiring diagram.

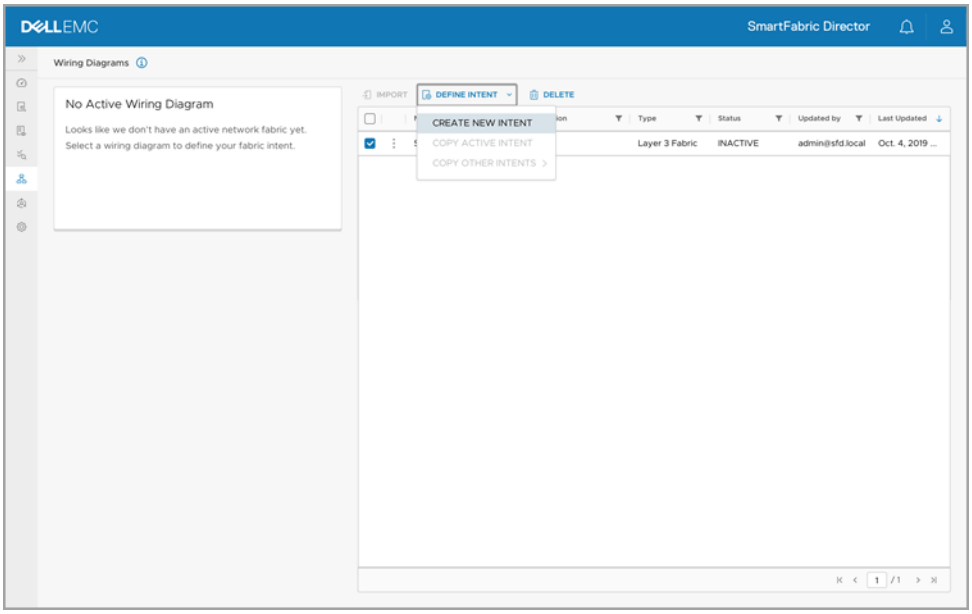

4. Select **View** to display the wiring diagram.

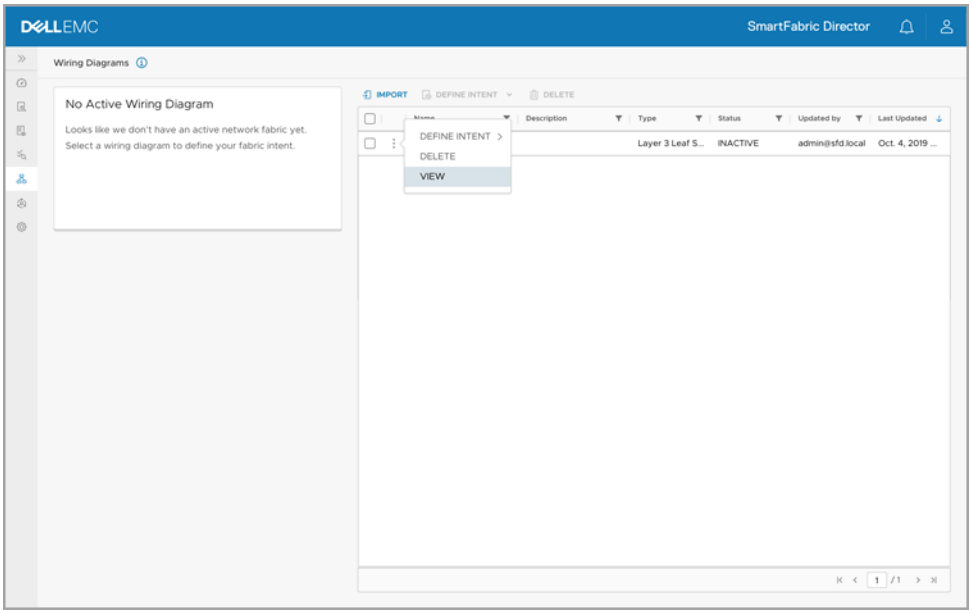

Imported wiring diagram.

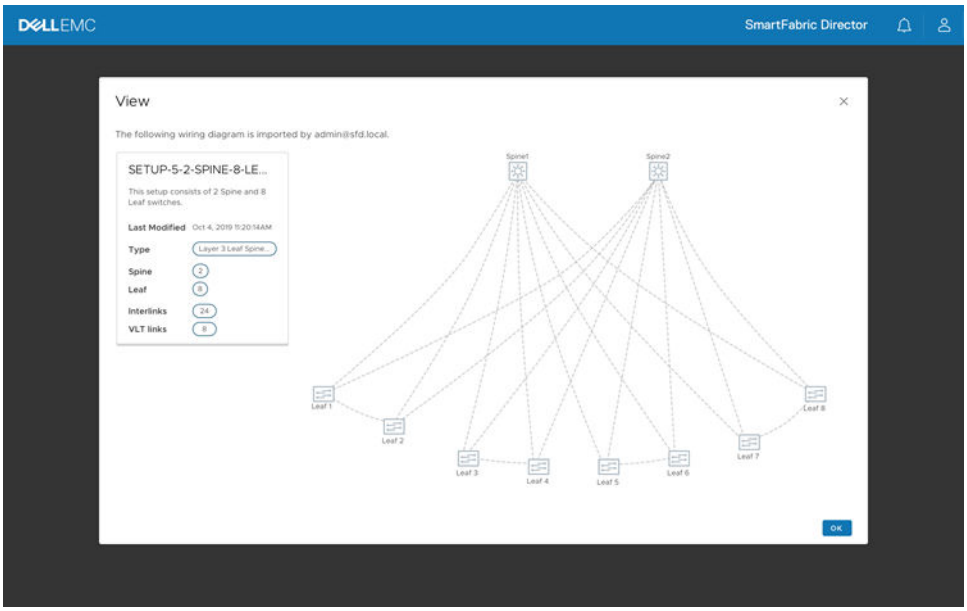

Define fabric intent

This information describes how to define a fabric intent. You can either import a previously defined fabric intent or start a new one from scratch.

Select fabric template

1. Click **Fabric Intent** from the left column to view and define fabric intents, then click **Get Started**. If there are no existing fabric intents, you can use the user interface to specify the fabric intent. Any such file can be used as a seed and edited.
2. Enter the **Fabric Name**, select a **Layer 3 BGP Leaf Spine Fabric**, then click **Next**.

   NOTE: Release 1.2.0 supports a Layer 3 BGP leaf spine fabric, or a Layer 2 leaf spine fabric.

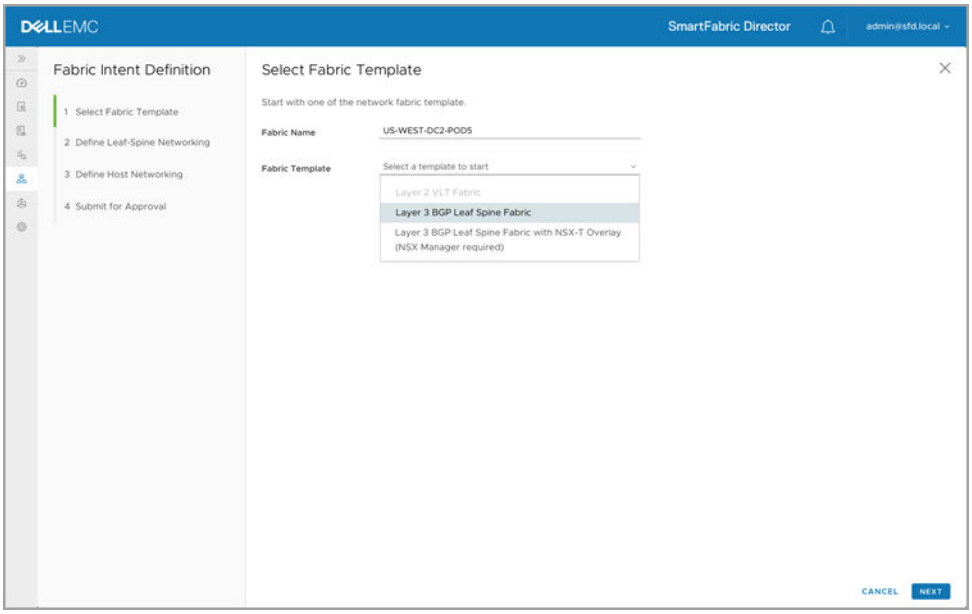

Layer 3 fabric

The leaf spine network template for a Layer 3 fabric is different than the Layer 2 fabric. This information explains the Layer 3 fabric.

Define leaf spine networking

You are now ready to specify the parameters to generate configuration for the interlinks between the leaf and spine switches.

1. For Layer 3 BGP leaf spine fabric, specify the leaf AS number range (start and end); spine AS number range (start and end), fabric interlink subnet, and the /32 loopback IP address seed. SFD generates a per-switch configuration for the interlinks between the leaf and spine switches.

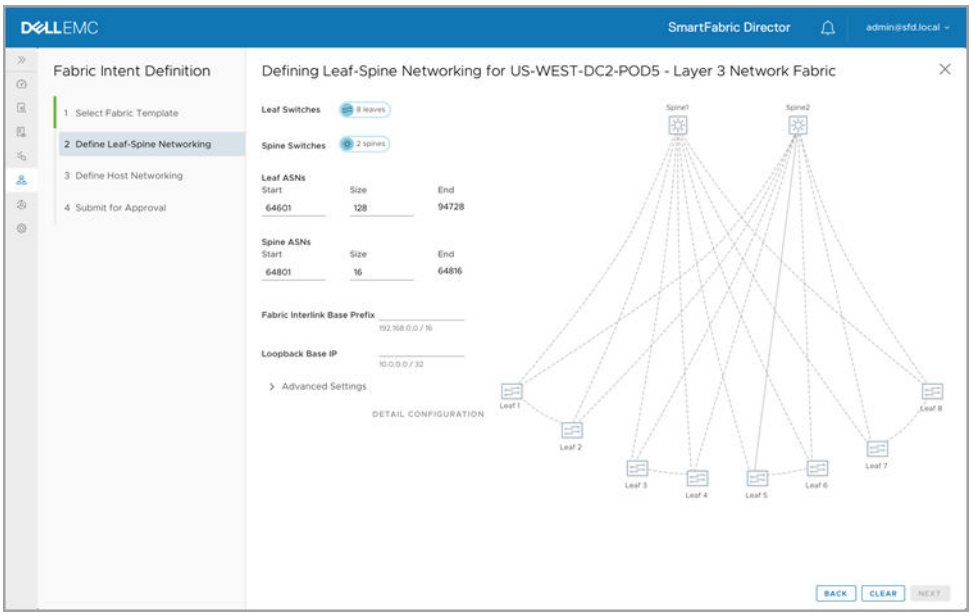

2. (Optional) Click **Detail Configuration** to view the interlinks between the leaf and spine switches, or click **Clear** to clear the specified intent, along with the detailed (per switch) configuration for the interlinks. Click **Back** to return, or click **Next** to continue.

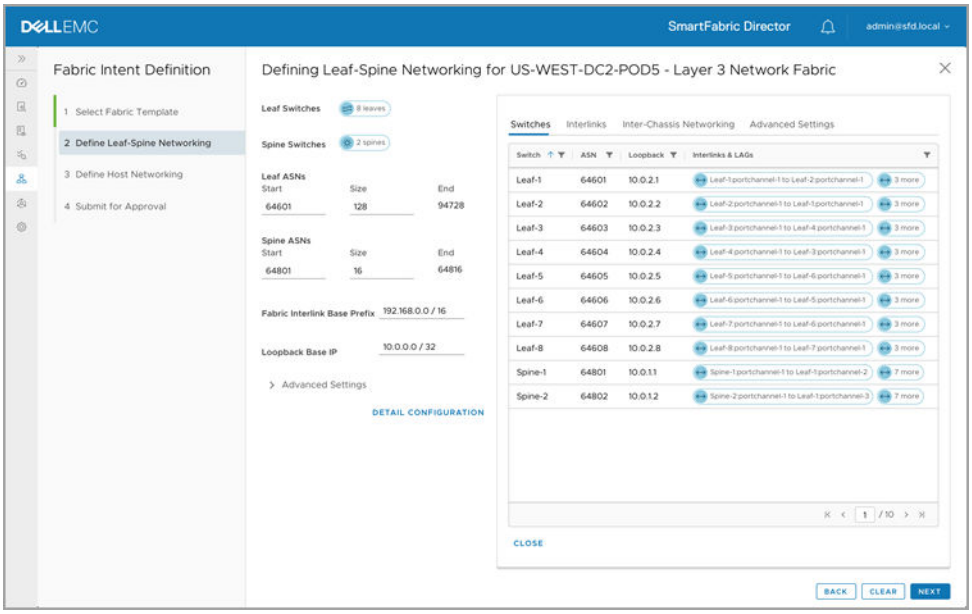

Details of interlinks configuration for Layer 3 fabric.

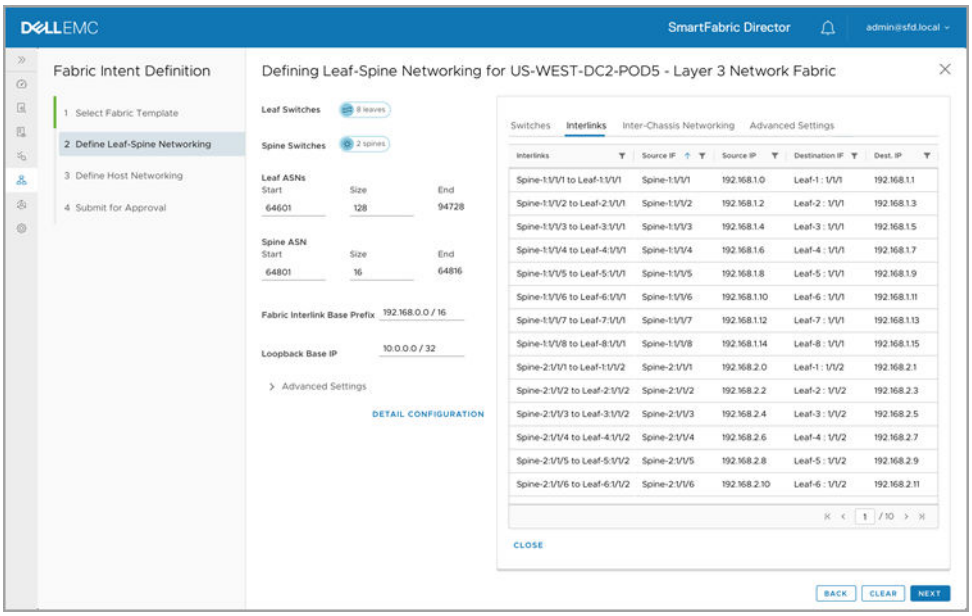

Details of interchassis configuration for Layer 3 fabric.

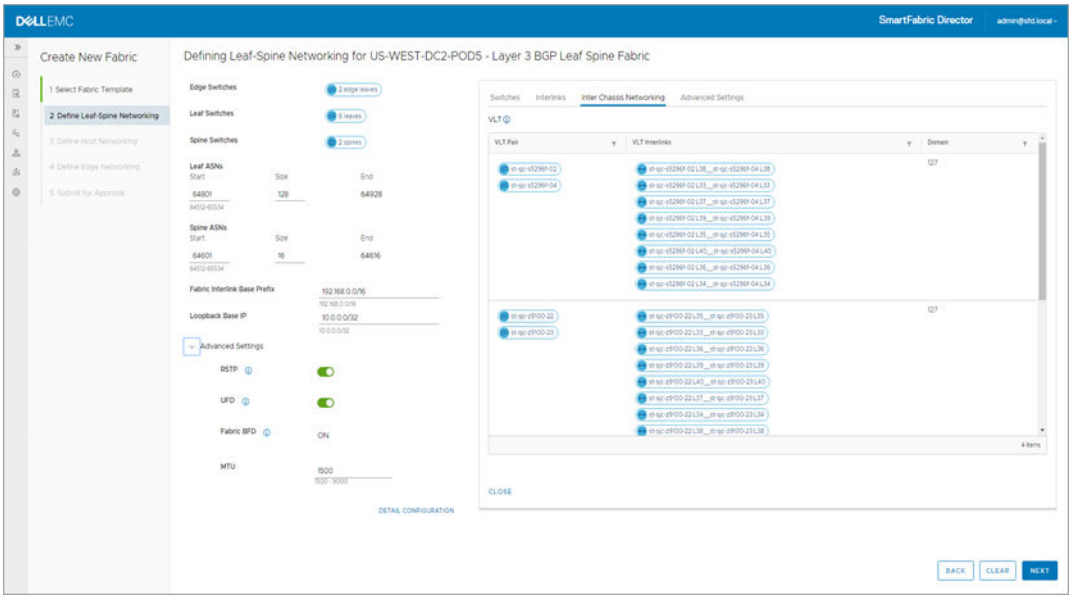

Details of advanced settings for Layer 3 fabric.

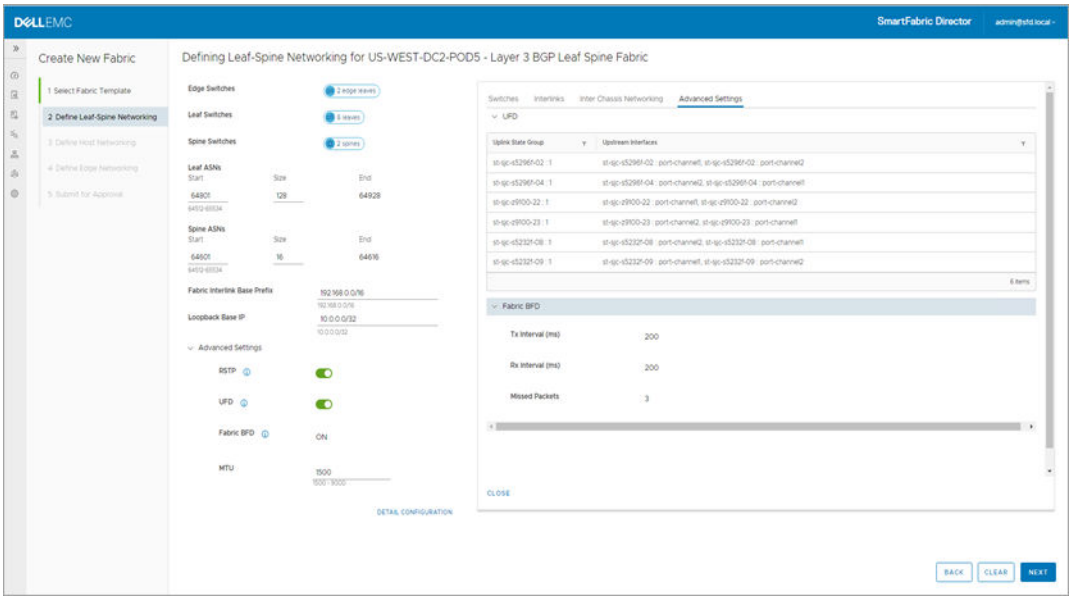

3. BFD is enabled by default on all interlinks between the leaf and spine, and can be configured to enable between edge leaf to external switch. In Layer 2 topology, BFD can be configured to enable between edge leaf and external switch. Select **Advanced Settings** to view the Fabric BFD settings.

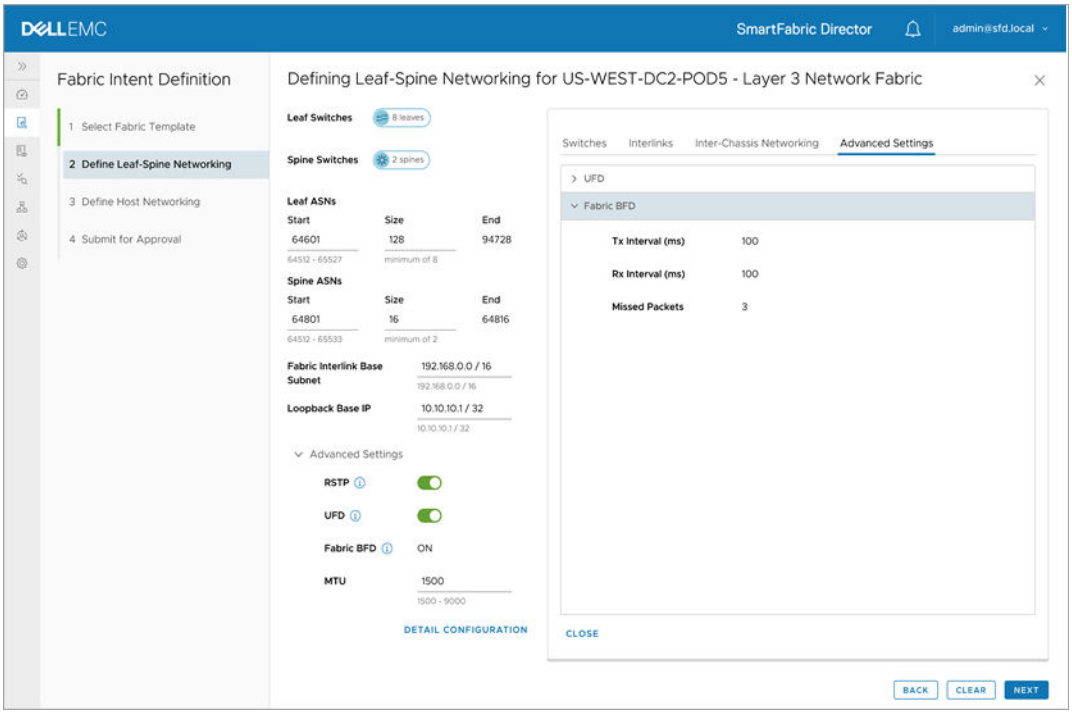

4. Specify the settings for SFD to connection with the vCenter; select **System Settings** from the left column, then select **VMware Manager Integration**. A list of vCenter Server connections previously configured display. If there are no existing vCenter Server connections, this table is blank.

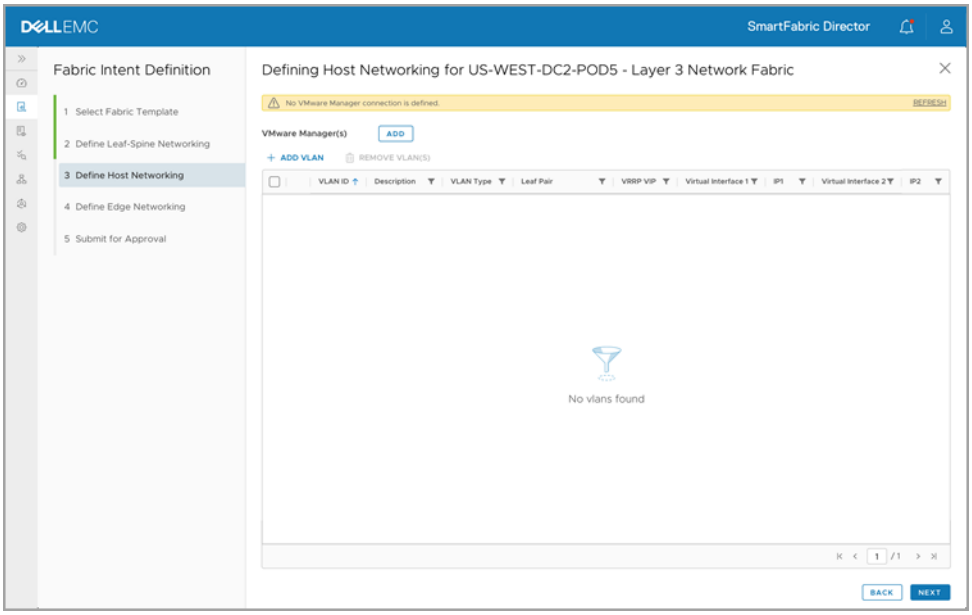

5. Click **Add VMware Manager**, select **vCenter Server**, enter the IP address/FQDN, enter the user credentials (username/password), enter an optional description (up to 255 characters), then click **Add**. The username/password does not display; click the eye icon to make the password visible.

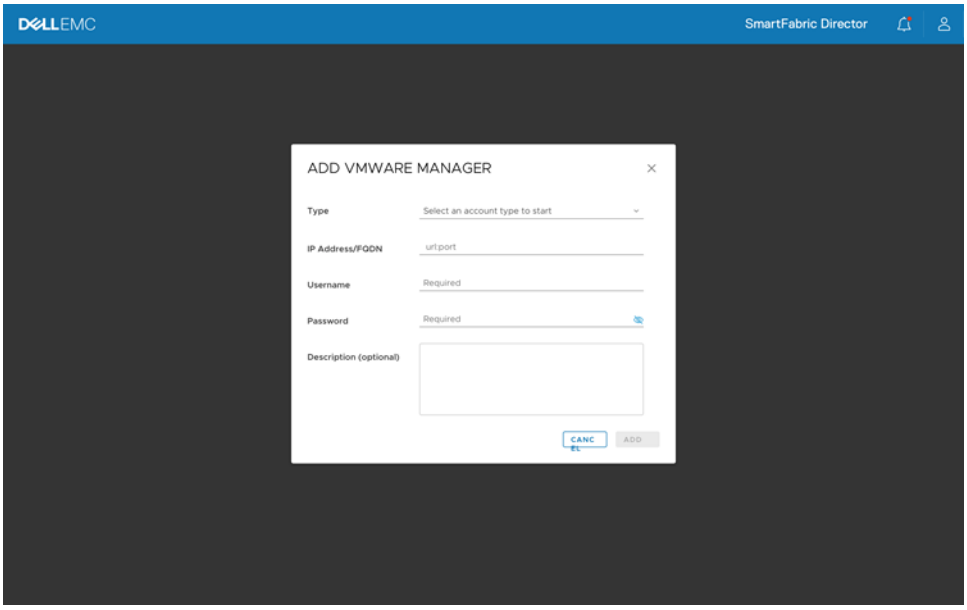

Define host networking

You are now ready to specify the parameters to generate configuration for host-facing ports, inter-VLAN routing, and host dual-homing. This screen also indicates the IP or DNS address of the vCenter (VMware Manager) used to manage the VMs on the hosts that are connected to this fabric.

1. Click **Add VLAN** to add the configuration for each port-group for the vCenter Server, enter the VLAN ID corresponding to a port-group, enter an optional description, then click **Next**.

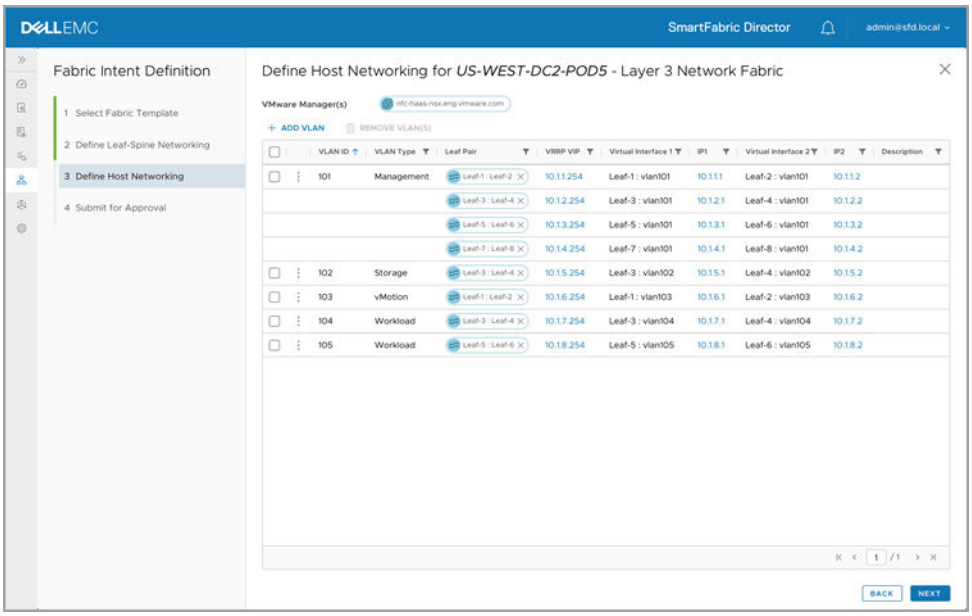

2. Associate one or more VLT pairs to a VLAN ID from the list of available VLT pairs (derived from the wiring diagram). Each leaf in the VLT pair has its own SVI IP, and each VLAN in the VLT pair has a VRRP virtual IP in the same subnet as the VLAN.
3. (Optional) Delete any VLAN ID row by selecting the checkbox to the left of each row.
4. Navigate between pages by using the arrows; click **Back** to go to the previous step, or click **Next**.

Layer 2 fabric

This information explains the Layer 2 fabric. The leaf spine networking template for a Layer 2 fabric is different from that for a Layer 3 fabric.

Define leaf spine networking

For Layer 2 leaf spine fabric, you do not need to specify any parameters. SFD generates a per-switch configuration for the interlinks between the leaf and spine switches.

1. Select **Configurations** from the left column to view and define fabric intents, then click **Get Started**. If there are no existing fabric intents, you can use the user interface to specify the fabric intent. Any such file can be used as a seed and edited.
2. Enter the **Fabric Name**, select **Layer 2 Leaf Spine Fabric**, then click **Next**.

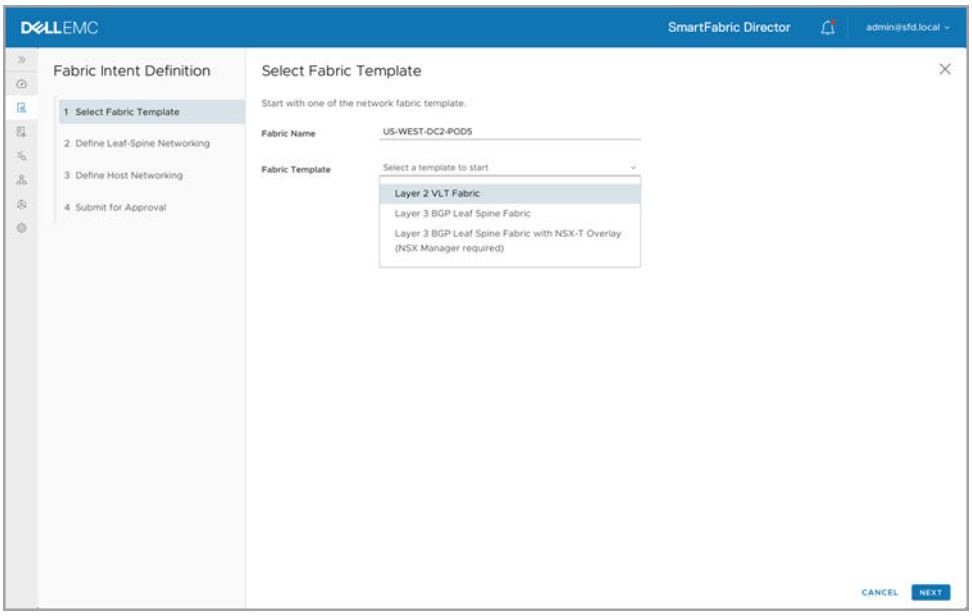

You are now ready to specify the parameters to generate configuration for the host facing ports, inter-VLAN routing, and host dual-homing. This screen also indicates the IP or DNS address of the vCenter (VMware Manager) used to manage the VMs on the hosts that are connected to this fabric.

1. For Layer 2 leaf spine fabric, SFD generates a per switch configuration for interlinks between the leaf and the spine switches. Click **Next** to continue.

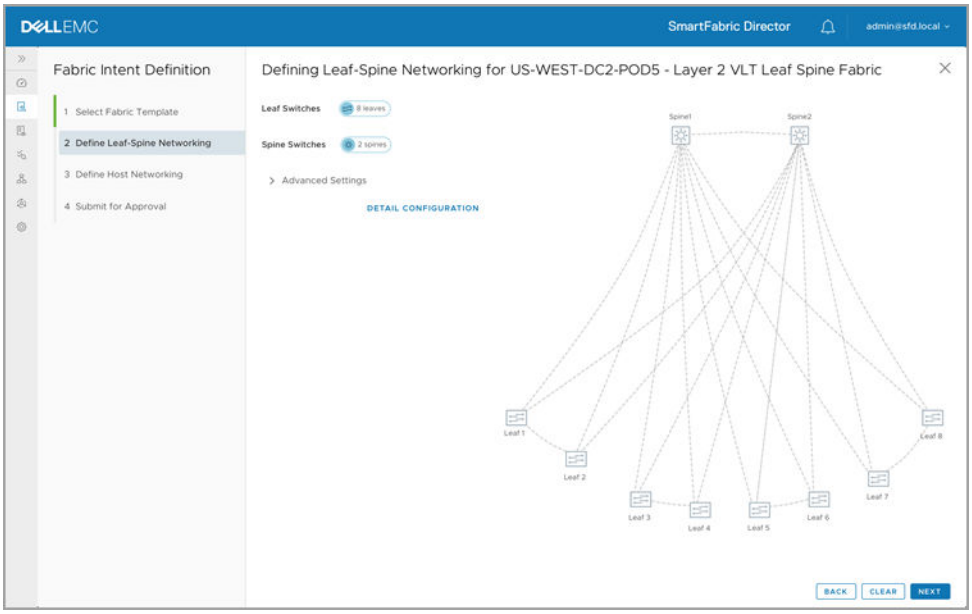

2. (Optional) Click **Detail Configuration** to view the per switch configuration, click **Clear** to clear the specified intent, along with the detailed (per switch) configuration for the interlinks. Click **Back** to return, or click **Next** to continue.

NOTE: Fields cannot be edited.

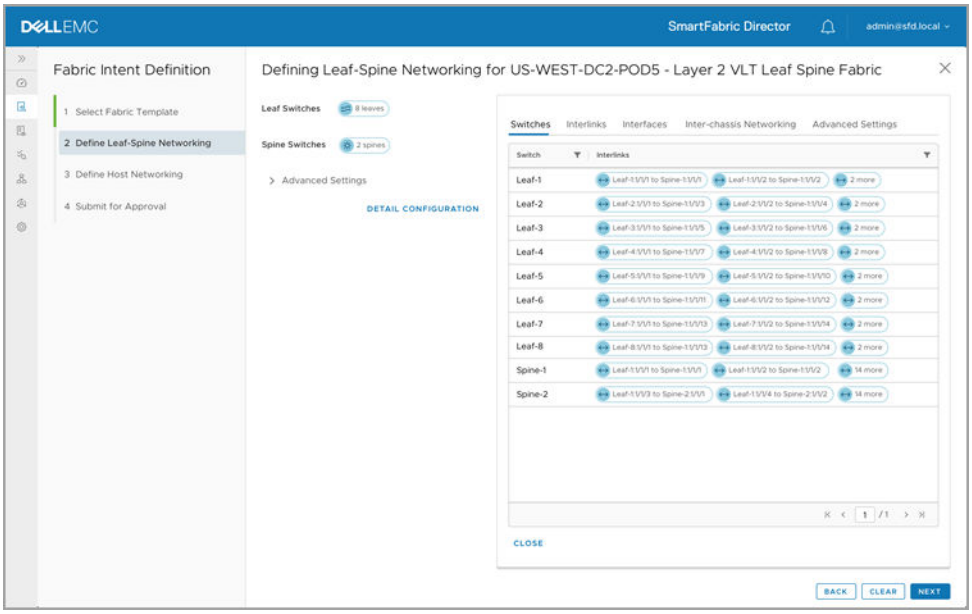

Details of interlinks configuration for Layer 2 fabric.

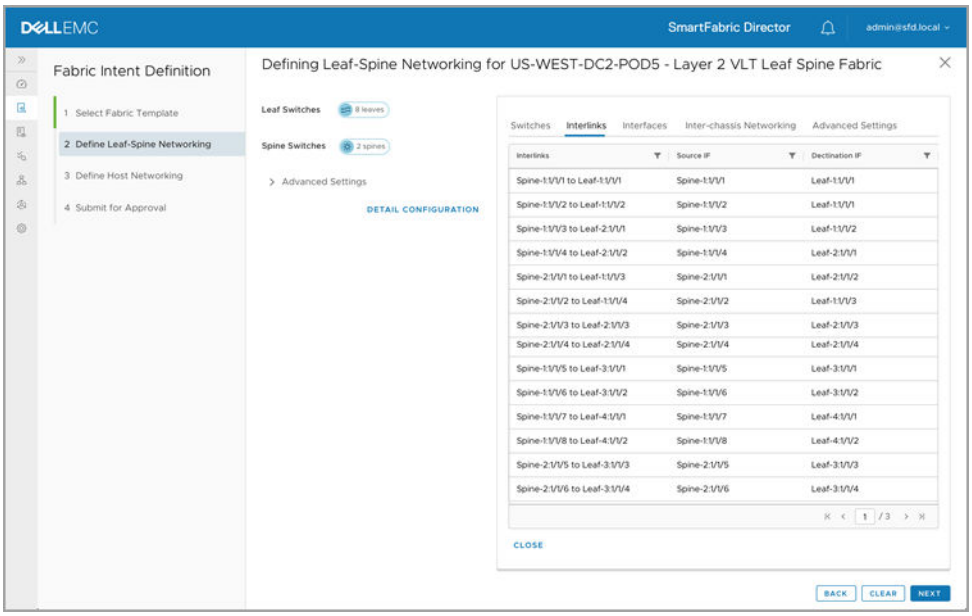

Details of interchassis configuration for Layer 2 fabric.

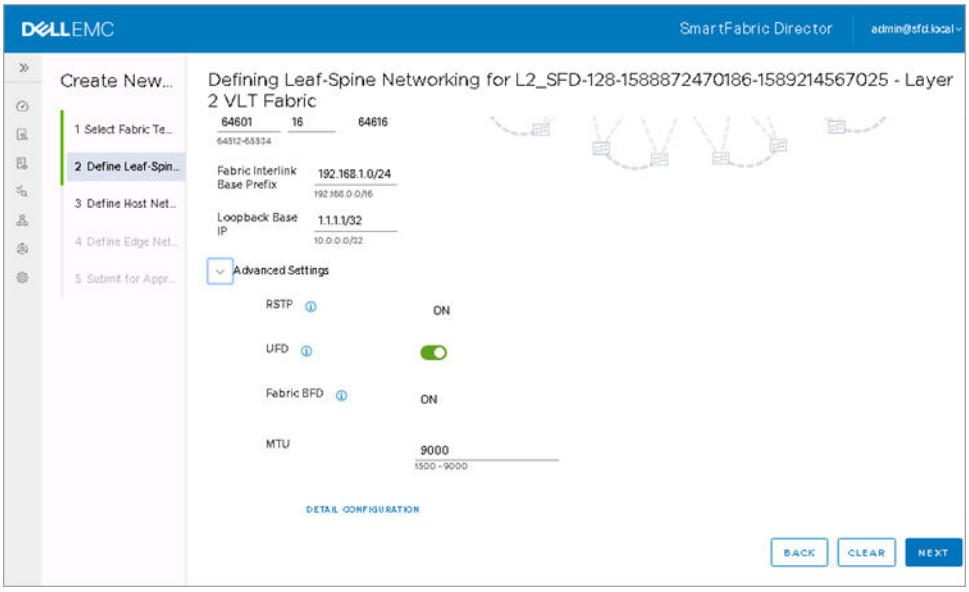

Details of advanced settings for Layer 2 fabric.

NOTE: RSTP is enabled by default and cannot be disabled.

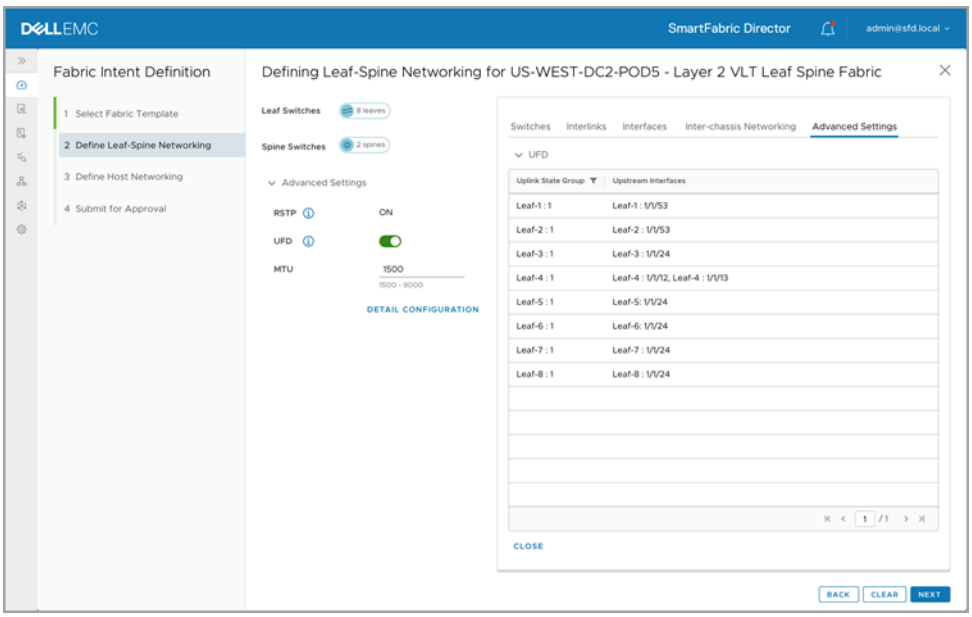

Define host networking

You are now ready to specify the parameters to be used by SFD to generate the configuration for host facing ports, inter-VLAN routing, and host dual-homing. This screen also indicates the IP or DNS address of the vCenter (VMware Manager) used to manage the VMs on the hosts that are connected to this fabric.

1. Select **System Settings** from the left column, then select **VMware Manager Integration**. A list of vCenter Server connections previously configured display. If there are no existing vCenter Server connections, this table is blank.

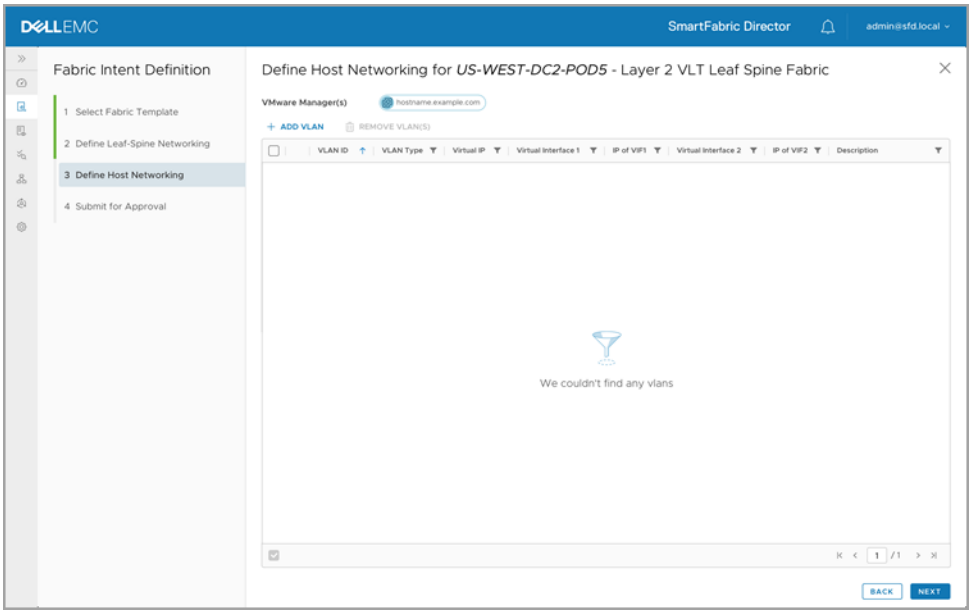

2. Click **Add VLAN** to add the configuration for each port group used by the vCenter, select the VLAN type, select the leaf pairs, enter an optional description, then click **Save**.

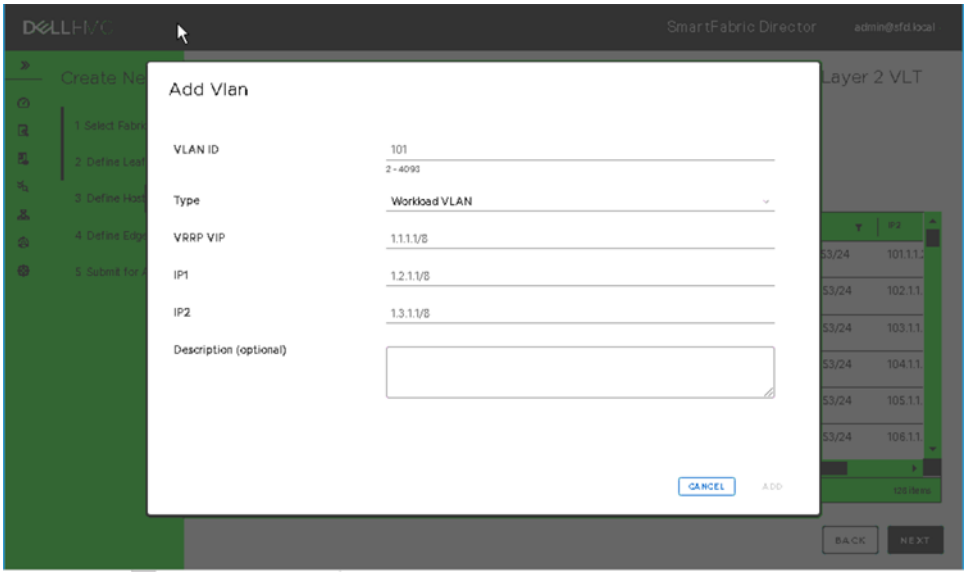

SFD creates a virtual interface on each spine, and associates it to the VLAN ID. Each virtual interface obtains an SVI IP, and each VLAN in the VLT pair obtains a VRRP virtual ip in the same subnet as the VLAN.

3. (Optional) Delete any VLAN ID row by selecting the checkbox to the left of each row.

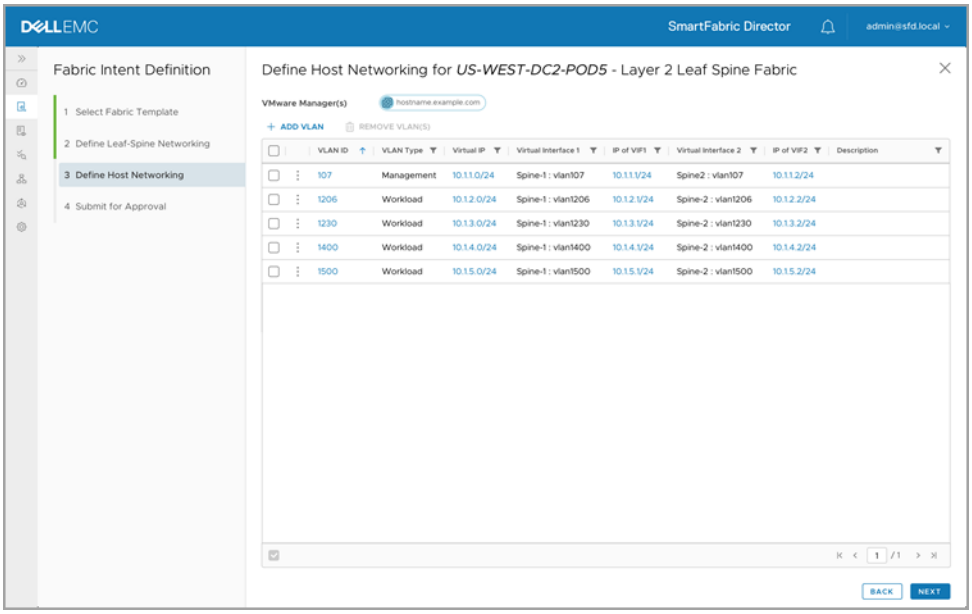

4. Navigate between pages by using the arrows; click **Back** to go to the previous step, or click **Next**.
5. Select **eBGP Settings** to enable BFD on the edge link, if needed.

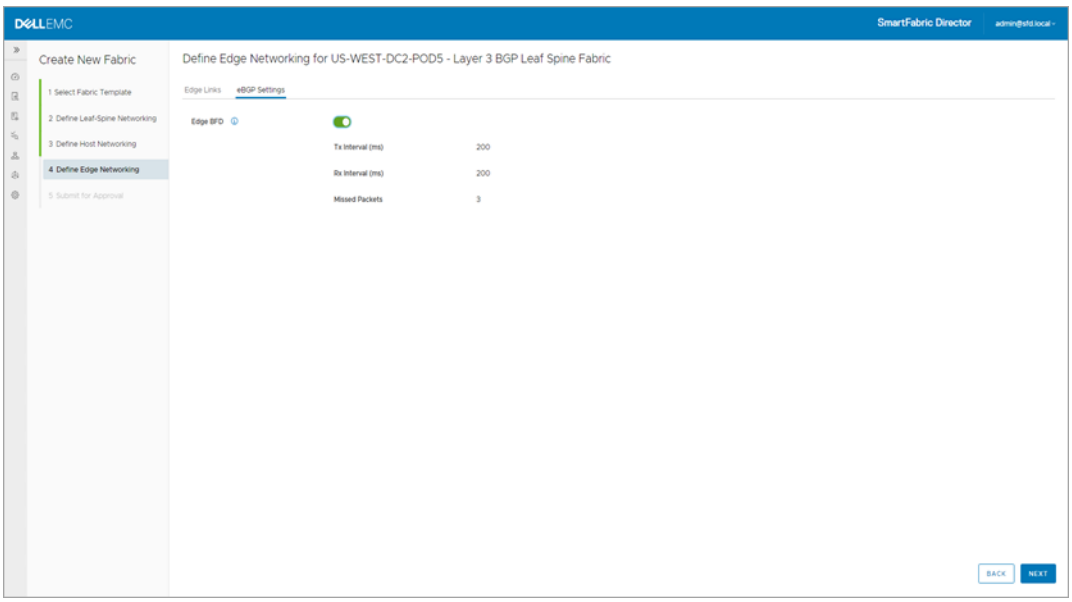

Define edge networking

Edge leaf switches provide external connectivity to the fabric. Edge leaf switches connect to upstream devices using a routing protocol to exchange reachability information. SFD supports the BGP routing protocol. The edge networking screens allow you to specify the required parameters to establish connectivity with an external device.

NOTE: The external device, typically a router, is not part of the fabric that is managed by SFD.

1. This optional step is required to configure the edge leaf networking. Click **Back** to return to the previous step.
2. Select **Edit** to configure an edge link.

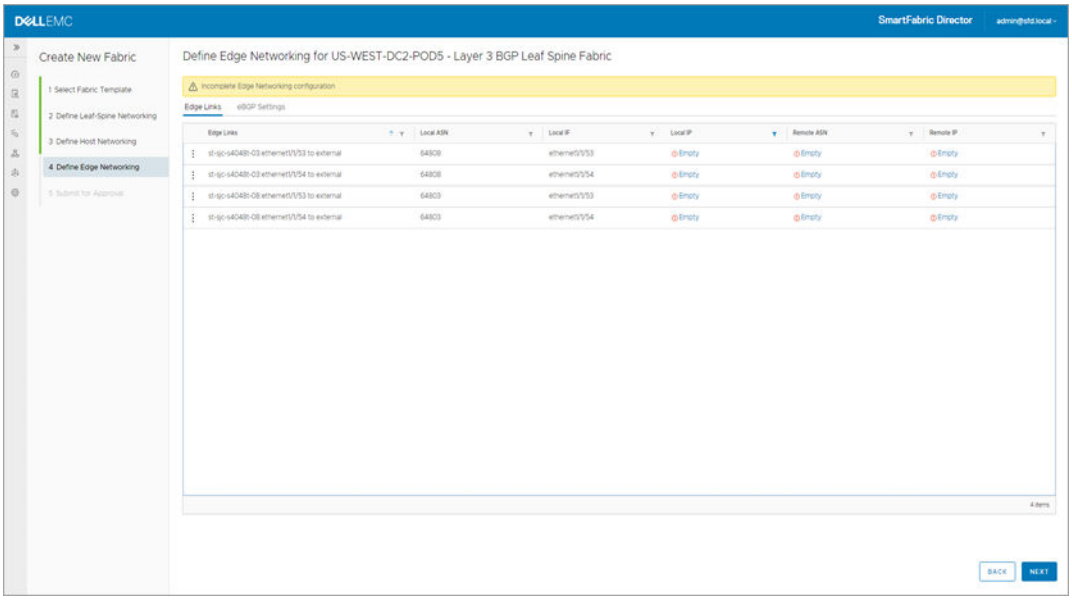

3. Enter the IP address for local, enter the remote ASN and IP address, then click **Save**.

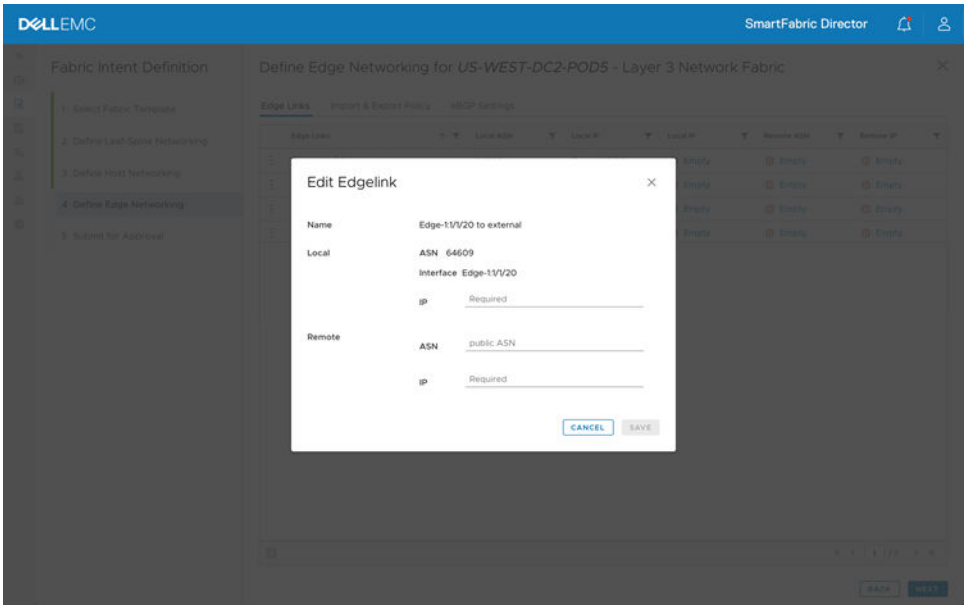

4. Continue editing edge links, then click **Next**.

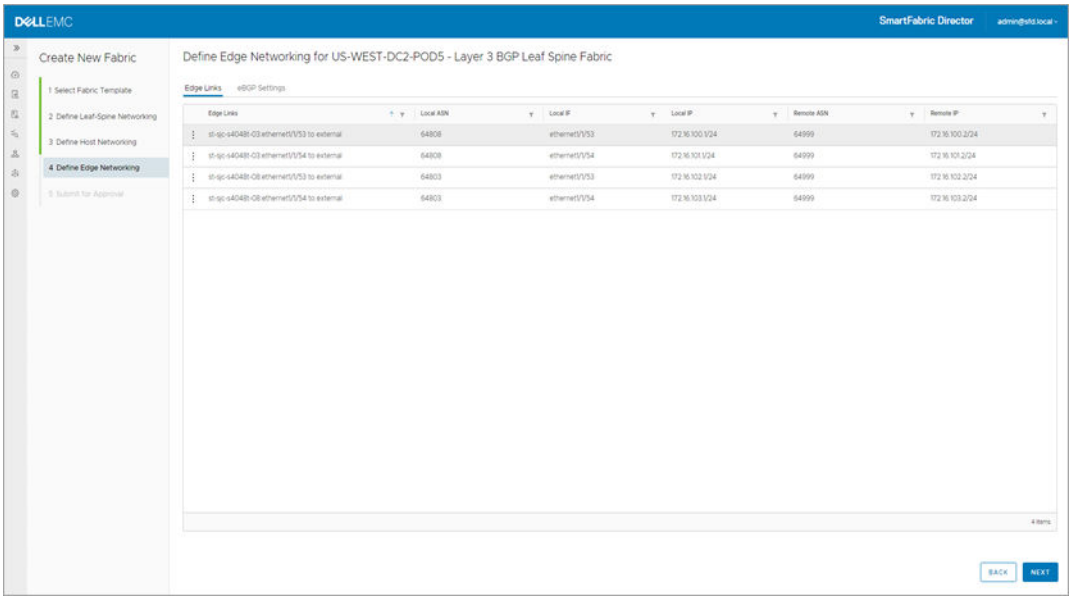

Define overlay networking

This optional step is required to configure the interface with NSX-T which performs the overlay switching, and the VLAN configuration to allow proper NSX-T operation. If the data center deploys overlay networking using NSX-T, SFD provides the ability to integrate with NSX-T, and install the necessary VLANs on switches in the fabric (underlay).

1. Click **Add** VMware Manager(s). The NSX-T Manager details can also be added in the VMware Manager Integrations tab under Settings and Administration.

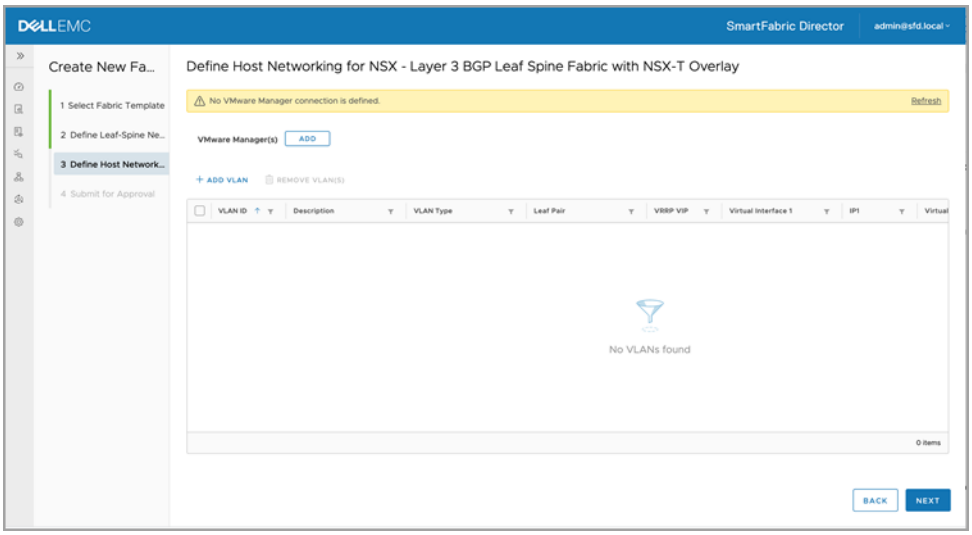

2. Select the **Type**, enter the **Username** and **Password**, enter an optional description, then click **Add**.

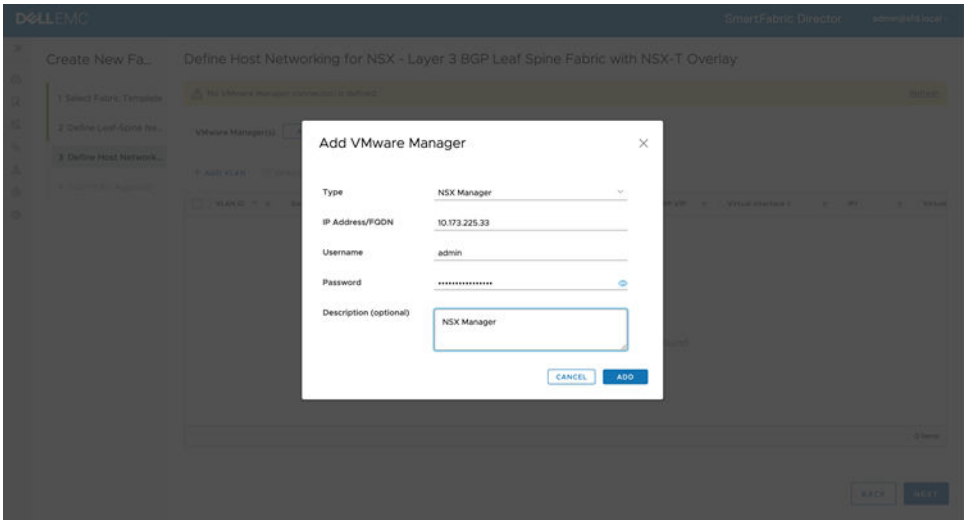

3. Select the VMware Manager, then click **Add VLAN**.
4. Enter the **VLAN ID**, then select the VLAN **Type**.

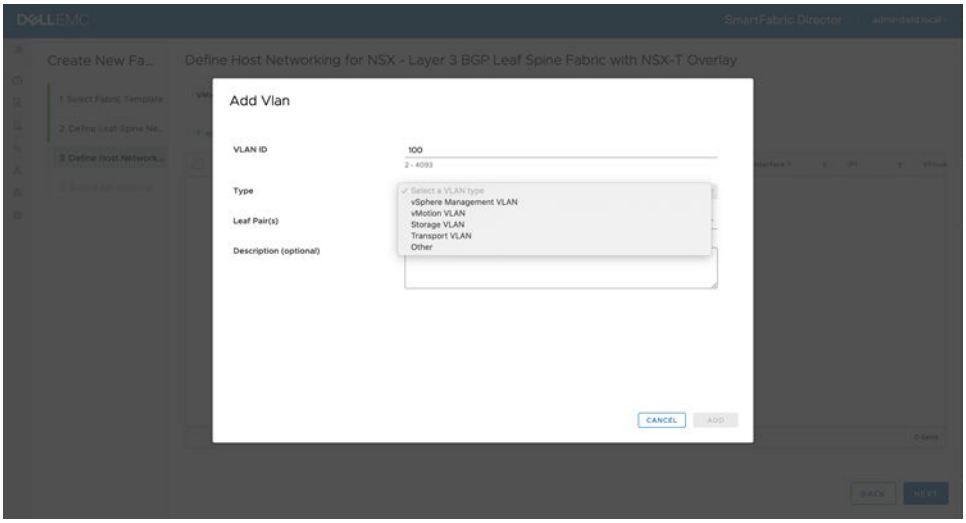

5. Assign **Leaf pair(s)** to each VLAN; repeat until you have assigned all leaf pair switches to the corresponding VLAN.

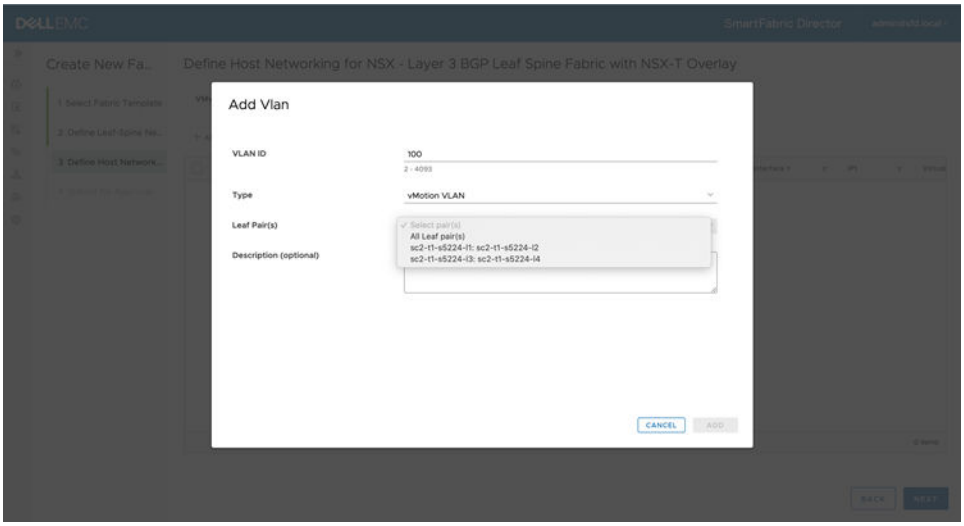

6. Specify the VRRP IP address and the IP addresses for the corresponding VLAN.

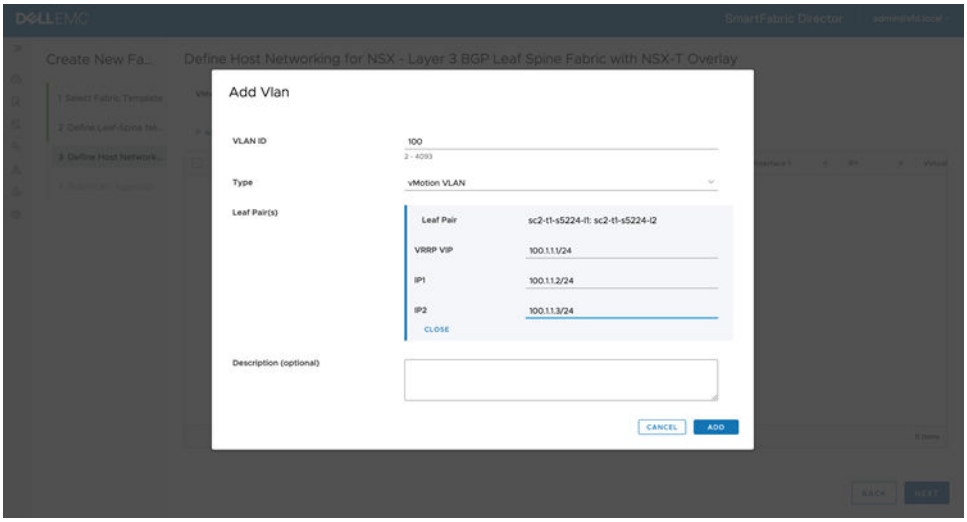

7. Enter an optional description, then click **Add**.

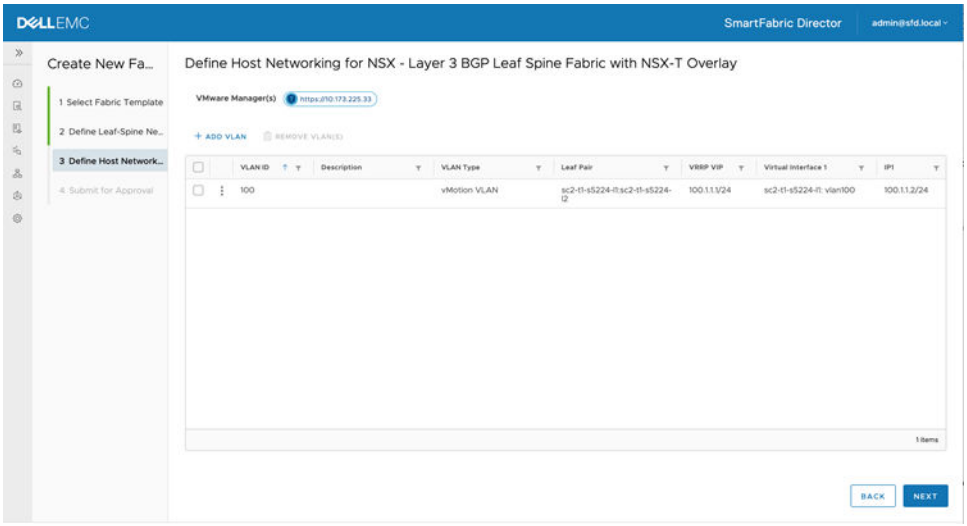

8. (Optional) Repeat the steps for all VLANs needed for NSX-T support.

Submit for approval

You are now ready to submit your L3 or L2 fabric intent for approval. The fabric intent must be approved before it can be deployed on the physical switches by SFD. Each fabric intent is associated with a wiring diagram.

1. The wiring diagram summary displays, along with a topology graph which corresponds to the wiring diagram. Click **Save for Later** to save the specified fabric intent as a draft in the SFD data store, or click **Back** to return to Define overlay networking.

   The summary displays different depending on the type of fabric configured. The example shows a Layer 3 fabric intent.

   NOTE: BFD is disabled by default on links from Edge ports to the external peer router. You can enable Edge ports if the external router has BFD enabled.

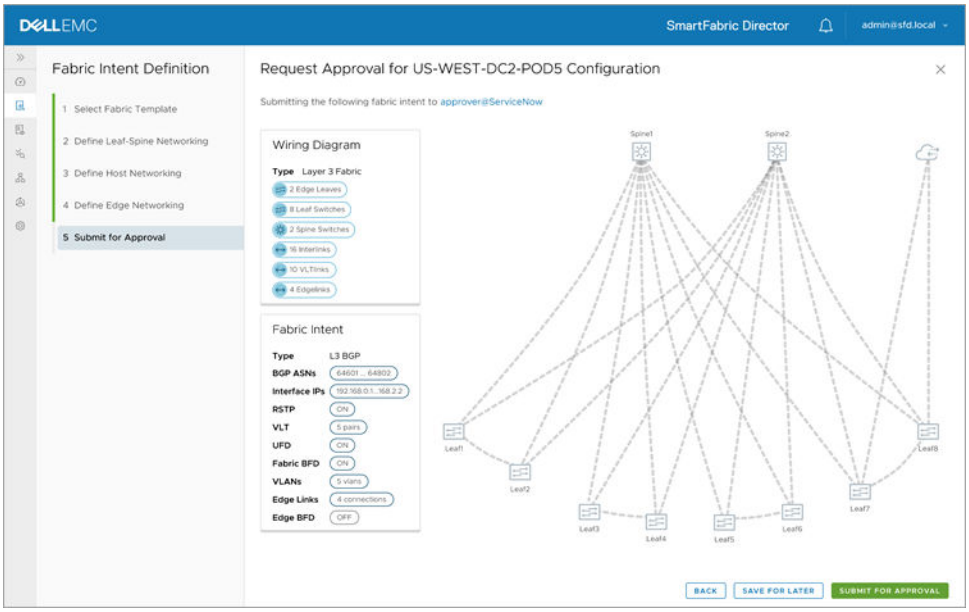

2. Click **Submit for Approval** to submit the fabric intent, along with the associated wiring diagram for approval by the authorized approver.
3. (Optional) Click **Back** to go to the previous step, or click **Save For Later** to save the specified fabric intent as a draft. All drafts are saved in the SFD data store.

Any fabric intent pending approvals are listed in the fabric intent list which can be viewed by selecting the Intent icon on the left.

Approve fabric intent

This information describes how to approve a fabric intent. Only Authorized Approvers can approve a Fabric Intent that is submitted for approval.

NOTE: In release of SFD 1.1.0, any user can approve an intent and the Fabric Intent is autoapproved and ready to deploy.

1. Select the checkbox to the left of the approved Fabric Intent to view details.

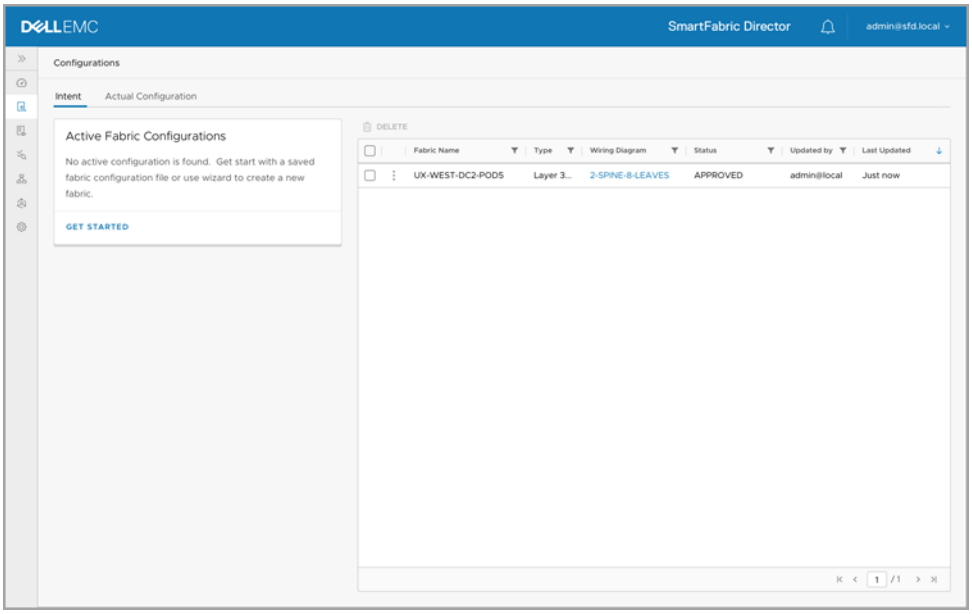

Deploy fabric intent

This information describes how to deploy the fabric intent. Deploying a Fabric Intent triggers configuration of individual switches of the fabric.

NOTE: You can only deploy Fabric Intents that have an Approved status.

1. The Deploy screen displays including a summary of the wiring diagram and the Fabric Intent (for reference), along with the topology graph which corresponds to the wiring diagram. The topology graph is updated based on the switches and links that are discovered by SFD. If one or more switches or links are either not discovered or are down (unreachable), you must review the summary and decide if you would like to deploy the Fabric Intent. Review the topology fully, then click **Deploy**.

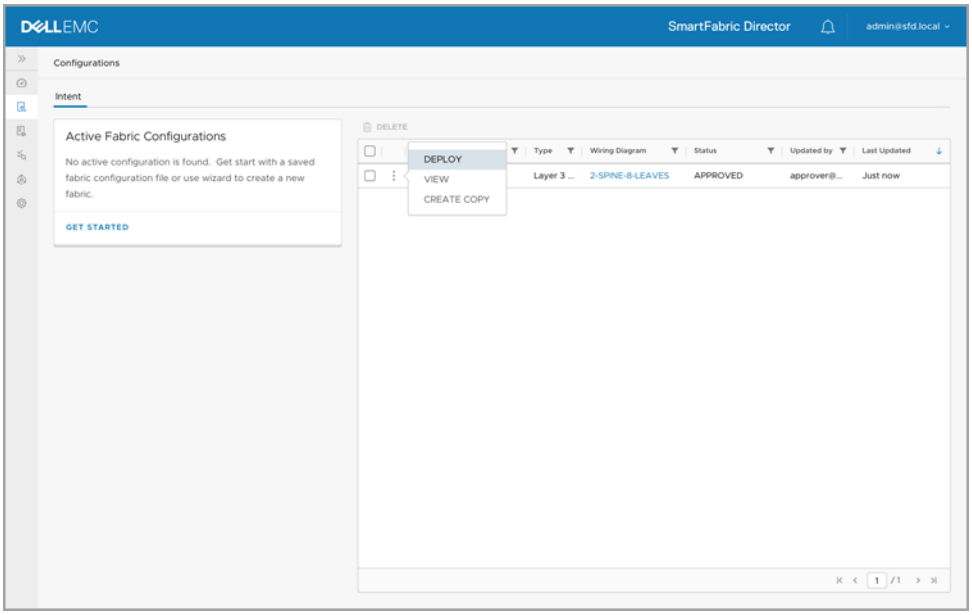

SFD starts configuring each discovered (and reachable) switch with the wanted configuration that is derived from the chosen approved Intent. SFD interfaces with each switch through the gNMI protocol for configuration, and the switches are configured simultaneously.

The SFD dashboard shows the progress of deployment using the progress bar located near the top of the screen. The switch in the topology graph turns green when configuration is successful. If all switches specified in the wiring diagram are discovered, reachable, and configured (and also the discovered topology matches the fully to the wiring diagram), the Fabric is deployed.

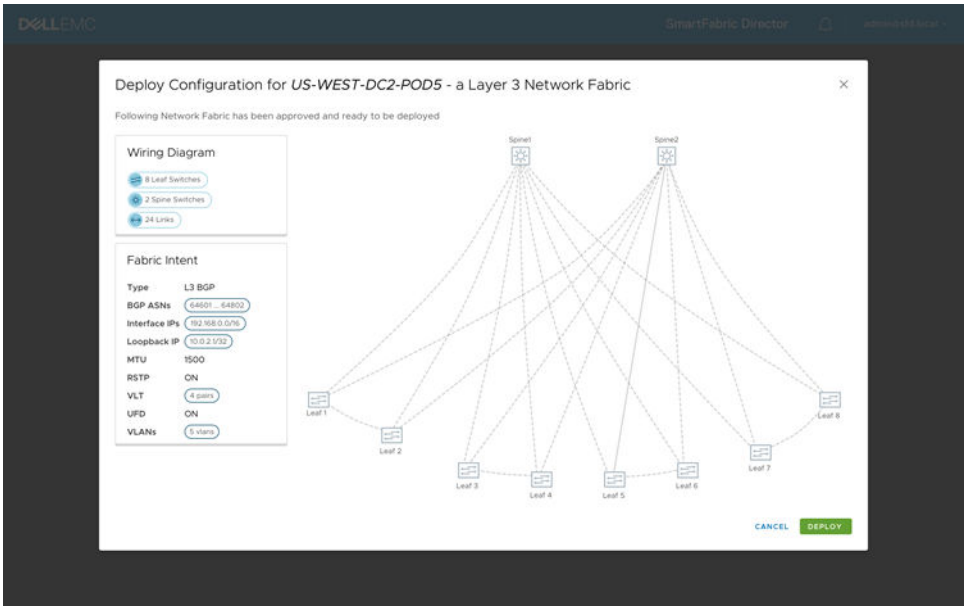

2. (Optional) Click **Cancel** to exit if you choose not to deploy the topology and return to the previous screen.

Add newly discovered switch to fabric

If a switch comes online or is discovered later and becomes reachable, SFD starts configuring the Fabric based on the wanted configuration—generated based on the deployed Intent.

The switch status is updated in the Topology graph but may not have more devices and/or links which are not specified in the Wiring Diagram (see Define fabric intent). A switch may discover a device such as an LLDP neighbor when that device is connected to an interface that is enabled.

Discovered devices that display in the discovered Topology but are not specified in the Wiring Diagram will not be shown in SFD. You must add newly discovered devices to the Wiring Diagram before you deploy the Fabric Intent.

Select the device then choose Add to go to defining a fabric wiring diagram. If you update the wiring diagram, you may need an update of the fabric Intent and associated approval, followed by deployment.

Discovery process

Once the Fabric Intent deploys, SFD starts the Fabric Discovery process. As part of the discovery process, SFD starts establishing a gNMI session with all the switches specified, using the IP address and credentials specified in the active wiring diagram.

Once a gNMI session is established, it uses the Openconfig objects to obtain details about the switch (service tag, hardware version, software version, and so on), and updates the switch and topology information in SFD.

The LLDP process on each switch is enabled and the LLDP table from each switch is read and used by the SFD to build a topology graph. This graph is referred to as discovered graph. If the gNMI session fails, then the switch state is updated accordingly.

The SFD is periodically notified by the gNMI Agent on the switch with its LLDP table information. The state of the switch (node) and the links (edges) between the switches are constantly updated based on the latest LLDP table infomation that is obtained, and the Discovery process is continuous.

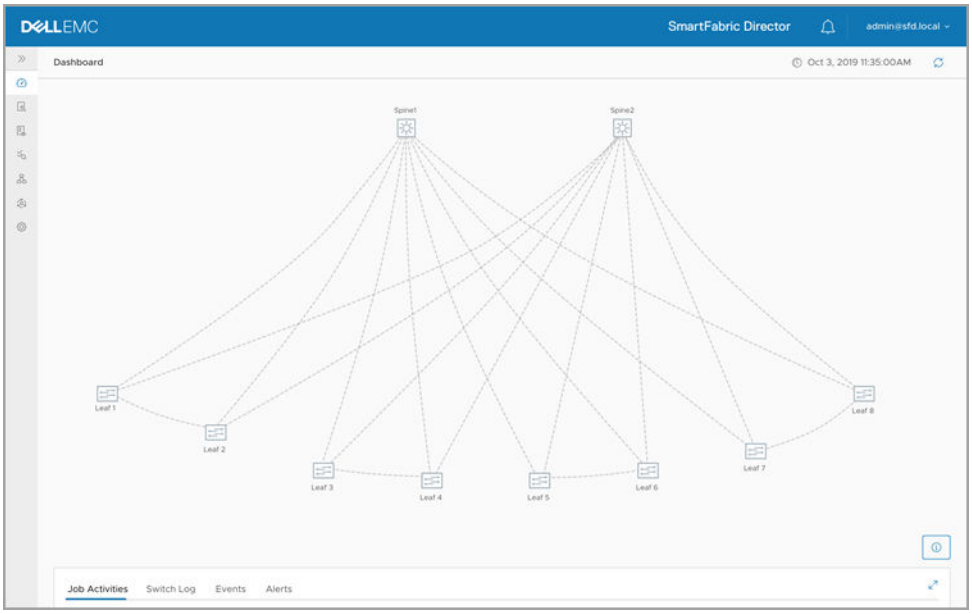

As the switches are discovered, SFD begins configuration of each switch. The progress can be viewed in the Job Activities pane at the bottom of the screen.

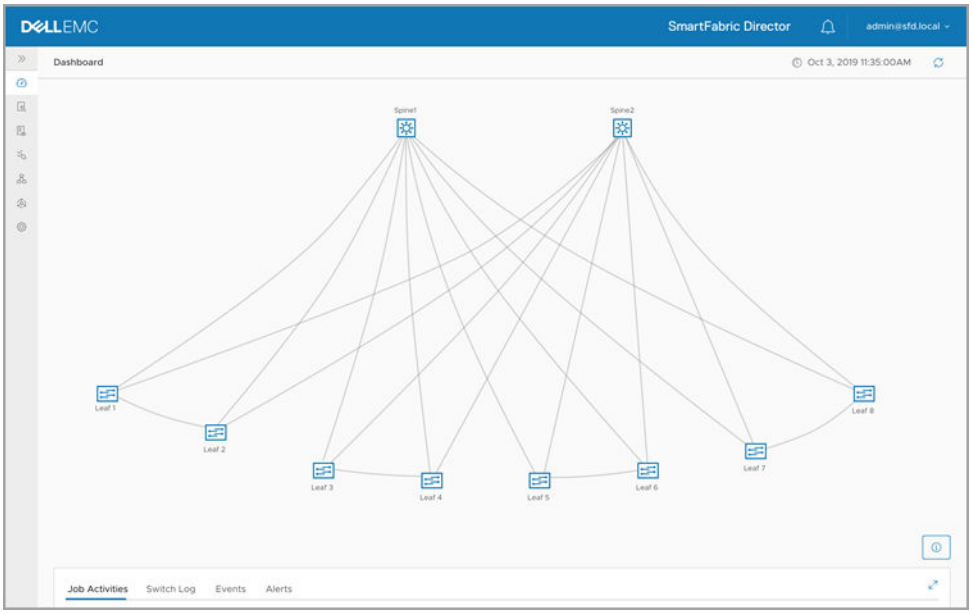

On successful configuration of a switch, the switch changes to green on the dashboard.

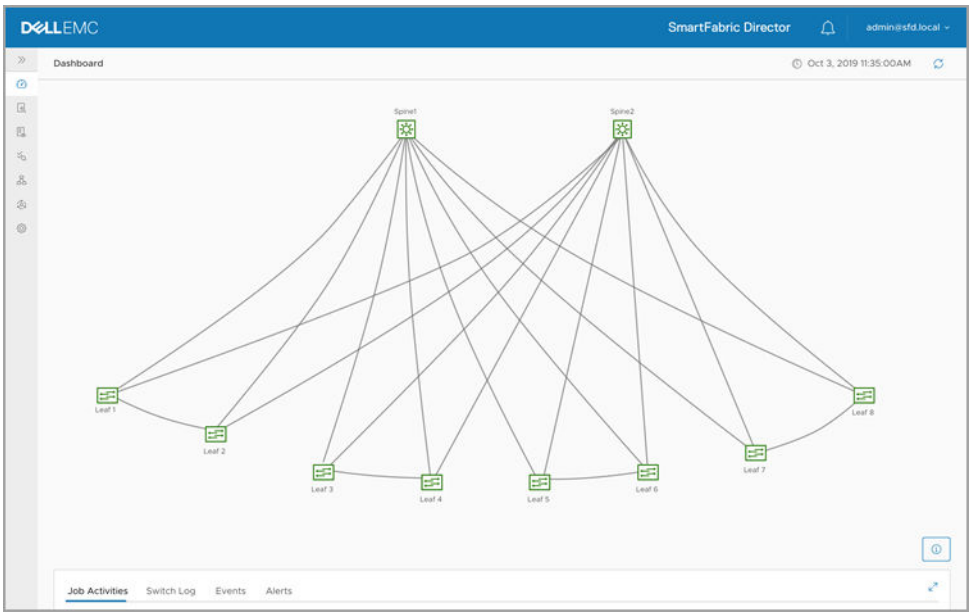

Once a fabric is deployed, the wiring diagrams table shows the corresponding deployed wiring diagram as Active.

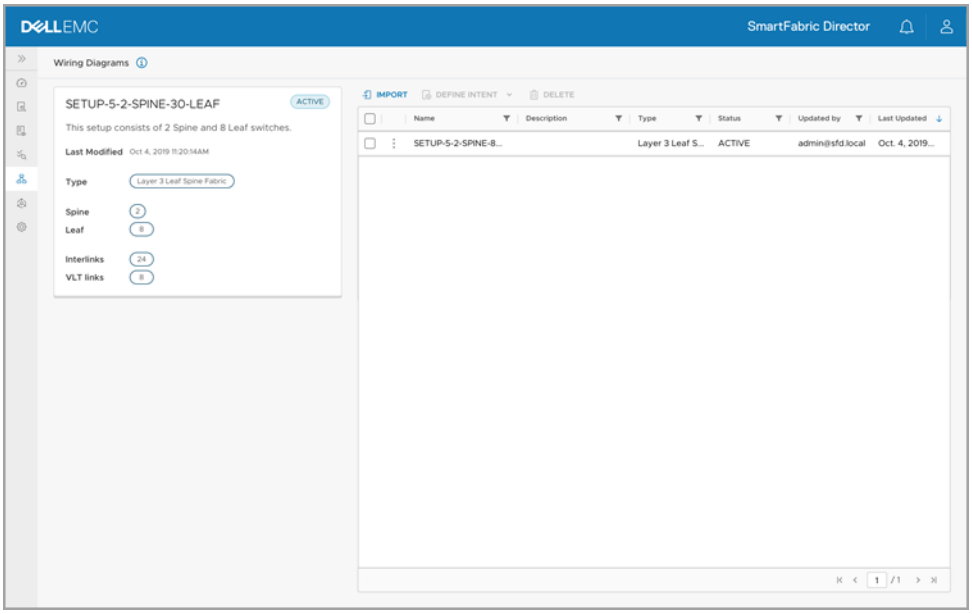

Reimport a wiring diagram

This information explains how to reimport a wiring diagram. You may update the fabric to add or remove switches, or add or remove links.

1. Create an updated JSON wiring diagram, then reimport it.

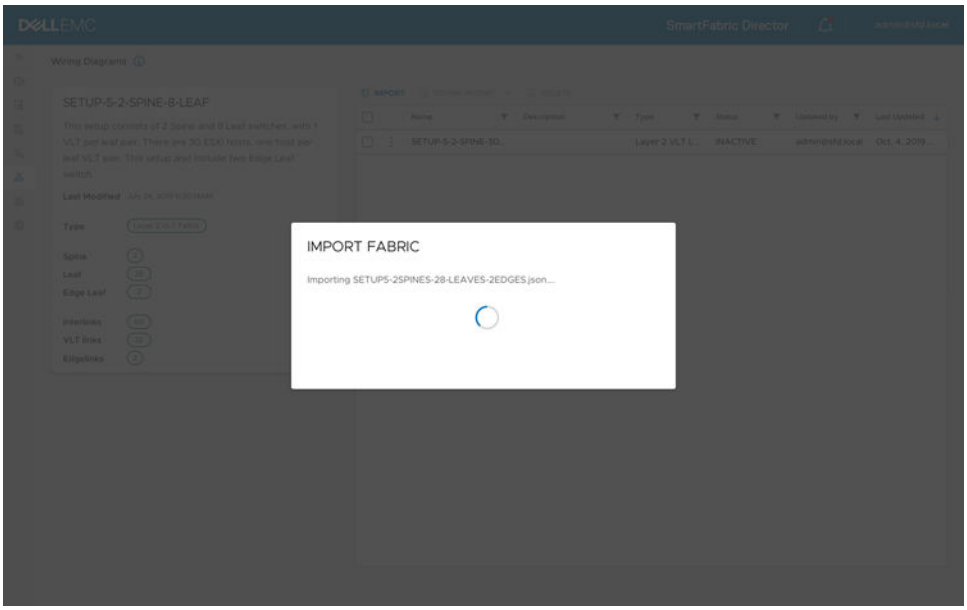

2. Define and associate a Fabric Intent. You can also clone or copy an existing Fabric Intent, then make any necessary changes by going through the Fabric Intent Wizard.

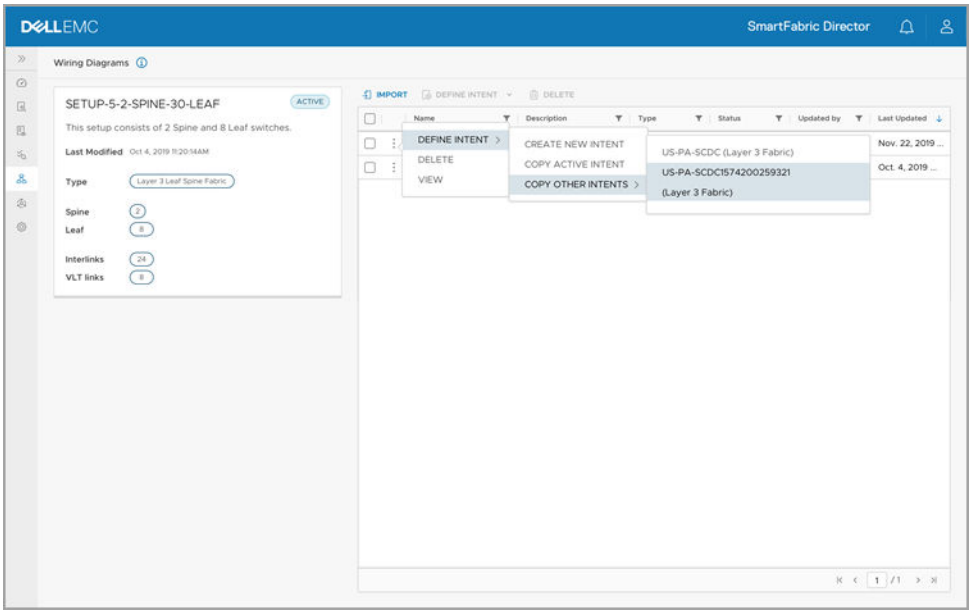

3. See Approve fabric Intent and Deploy fabric intent.

8

Monitoring

This information explains how to monitor the fabric, viewing switch level details. Monitoring data is the last information streamed by the switch to SFD. The switch streams telemetry information to SFD every 60 seconds. The streaming status of a network device is available in the switch profile under monitoring.

1. From the SmartFabric Director dashboard, select a switch to open up the switch profile panel to view a summary of the switch.

   NOTE: When viewing CPU data over a large time window, the graph displays peak CPU data. This time window can be adjusted or lowered to get a granular view of the CPU data.

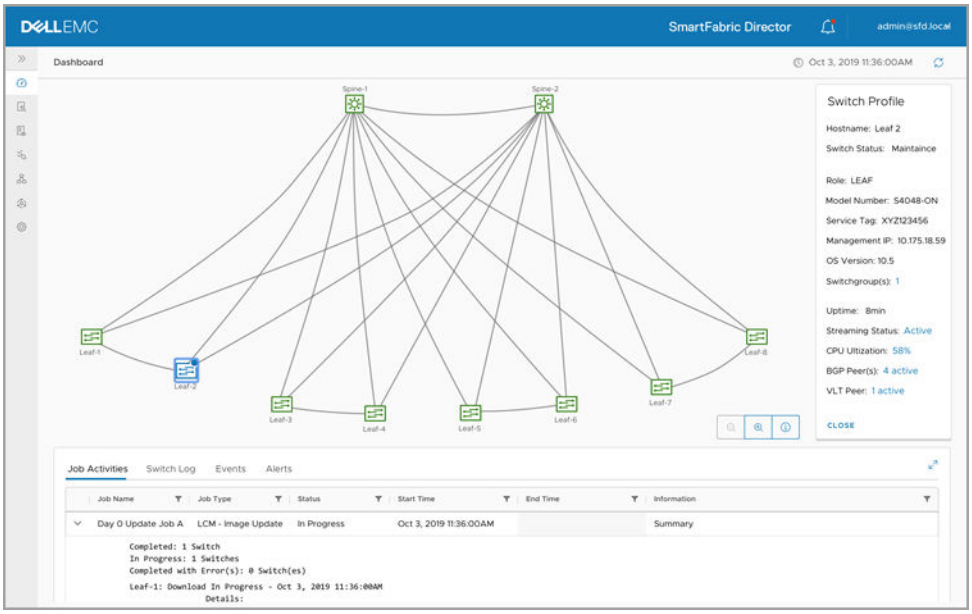

2. Select the **Monitoring** icon from the left to view the fabric health.

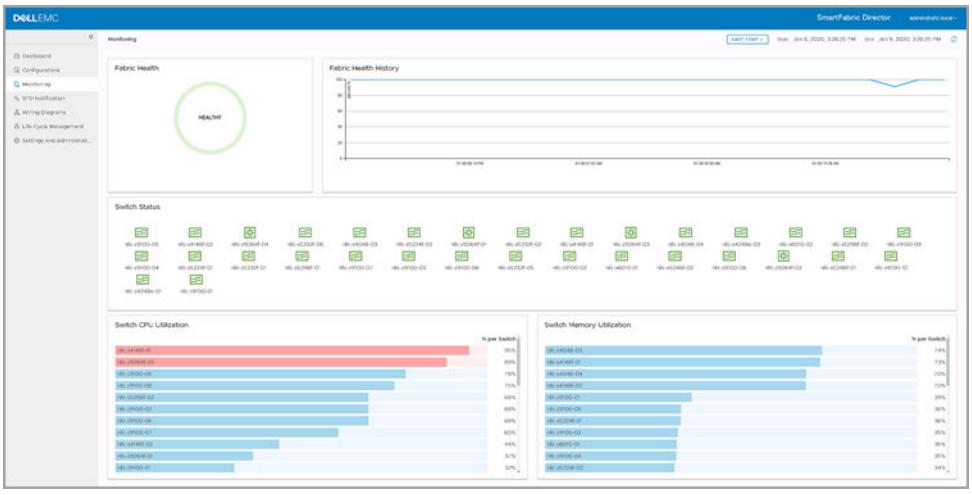

3. Select any interface to view the switch profile and details.

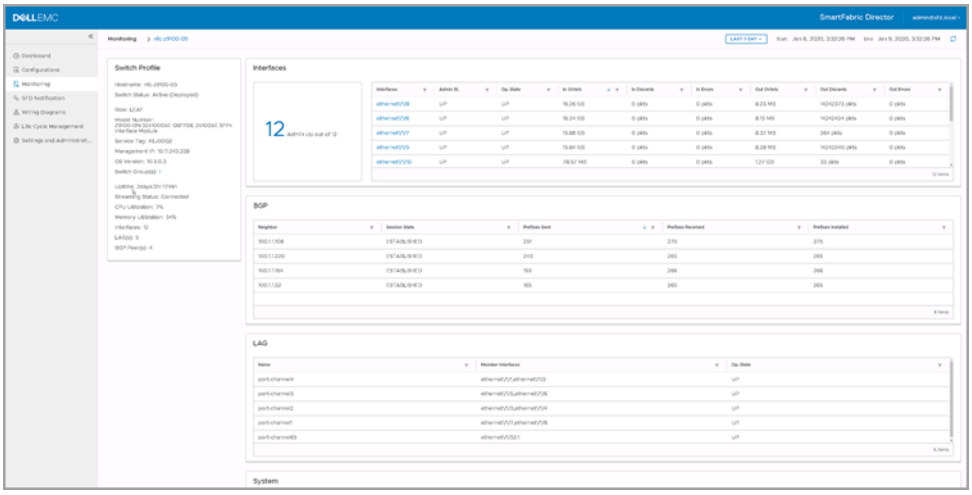

4. Scroll down to view additional details.

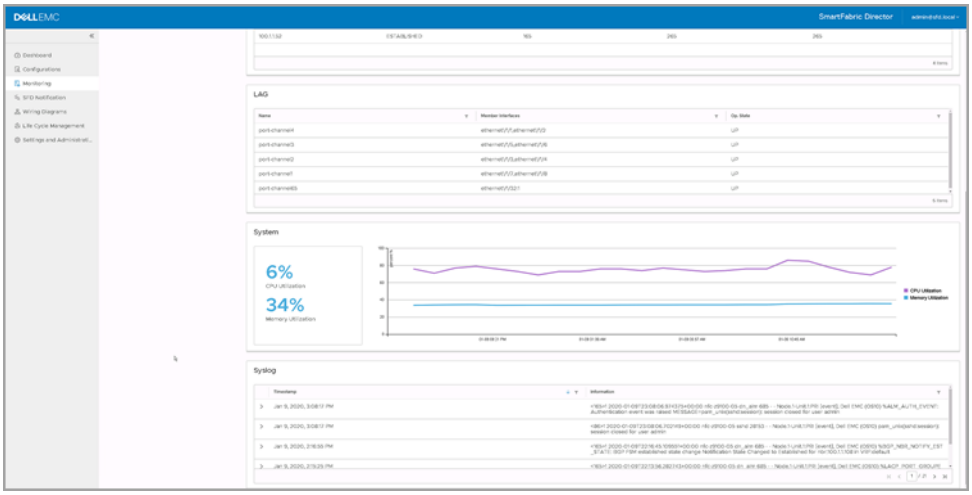

Topics:

- SFD notifications

SFD notifications

This information explains how to use the job activities, events, alerts, and the switch log to manage SmartFabric Director.

1. Select the **SFD notification** icon from the left to view job activities for fabric intent deployment.

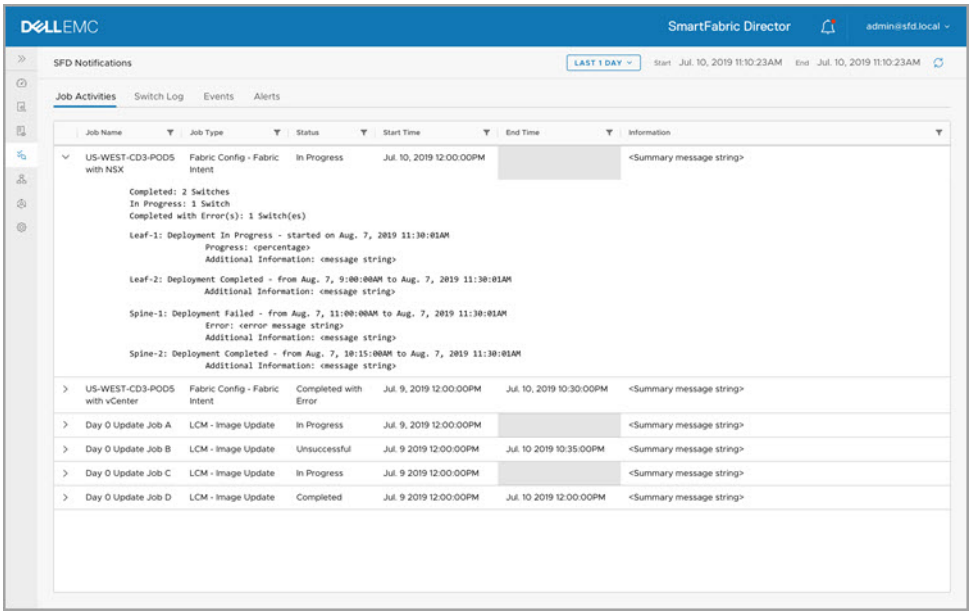

2. Select the **Switch Log** tab to view the system log information.

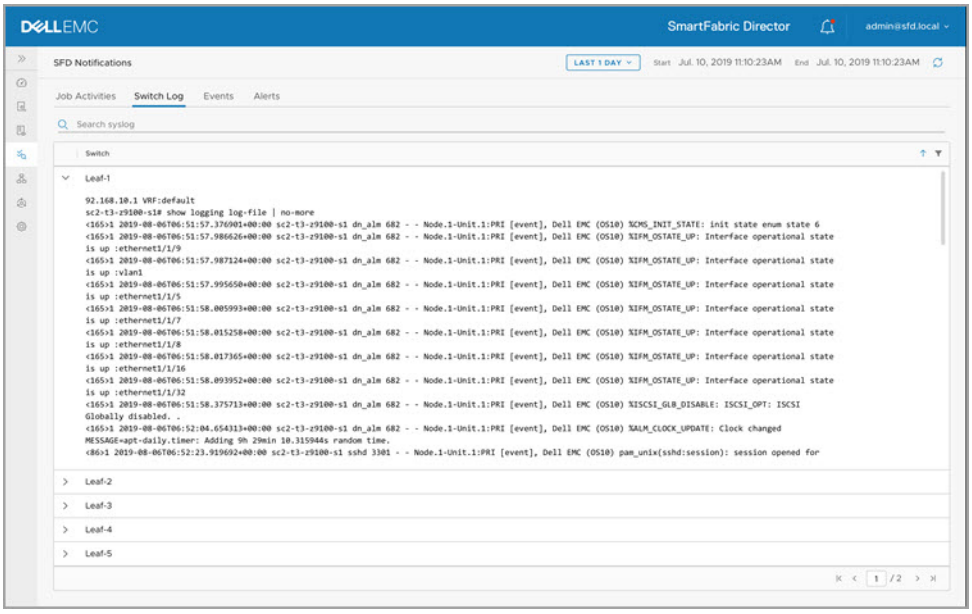

3. Select the **Events** tab to view all SFD events.

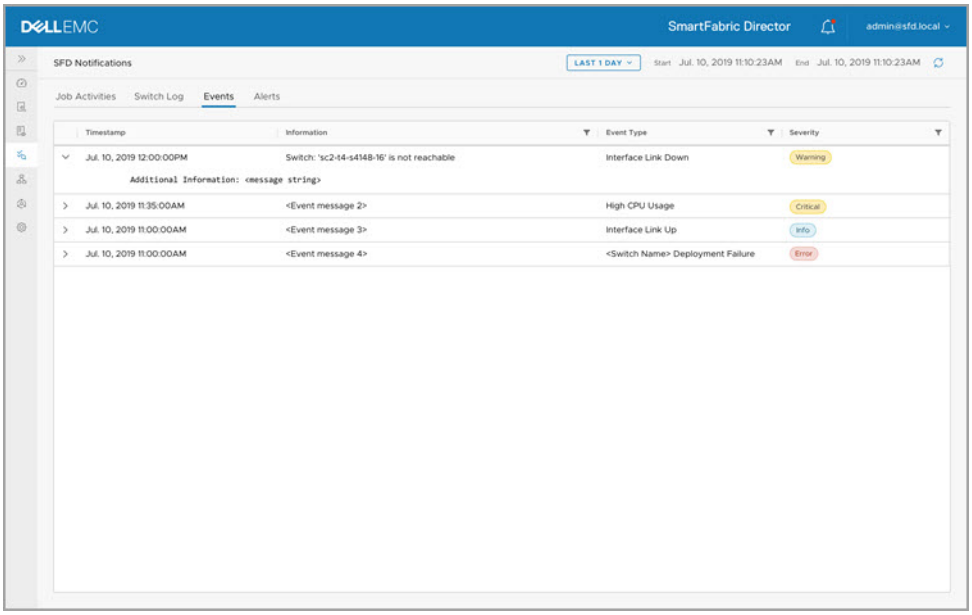

4. Select the **Alerts** tab to view all SFD alerts.

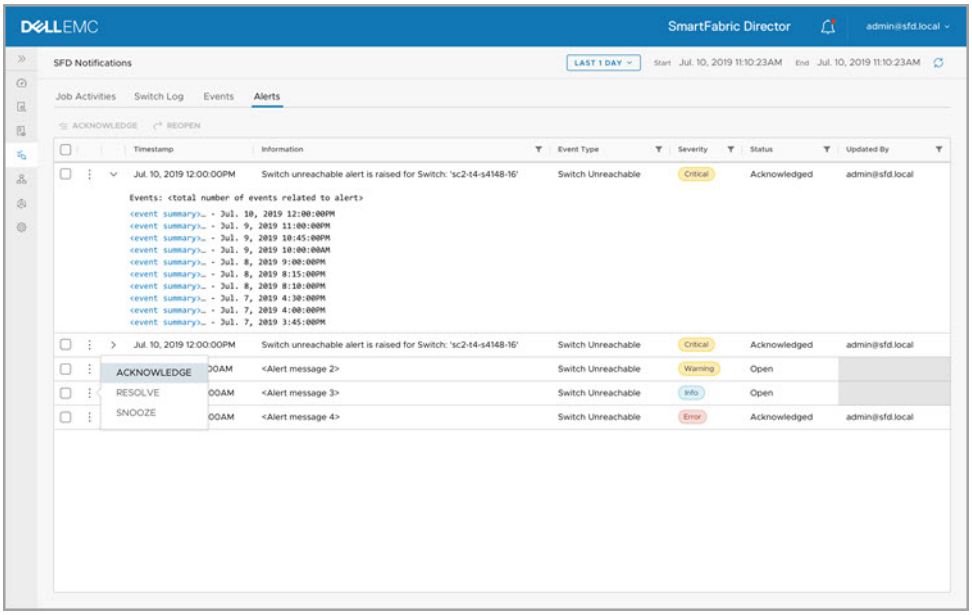

9

Switch lifecycle management

This information explains switch lifecycle management. Each switch in the data center fabric must have the same software image. You can upgrade or downgrade the switch image software using SmartFabric Director.

SFD is aware of all supported switch models and corresponding SmartFabric OS10 software images. This ensures that unsupported models and images are not deployed in the fabric.

Provide a file server (SFTP, FTP, SCP, TFTP, or HTTP) that is accessible through the Management port of switches, and reachable from SmartFabric Director. Download one or more relevant switch software images and manifest file to these servers.

Topics:

Specify switch image server

This information describes how to specify the Image Server where the switch software images are stored. See Download SFD image for more information about how to download a software image.

1. Select **Settings and Administration** > **Switch Image Servers**.

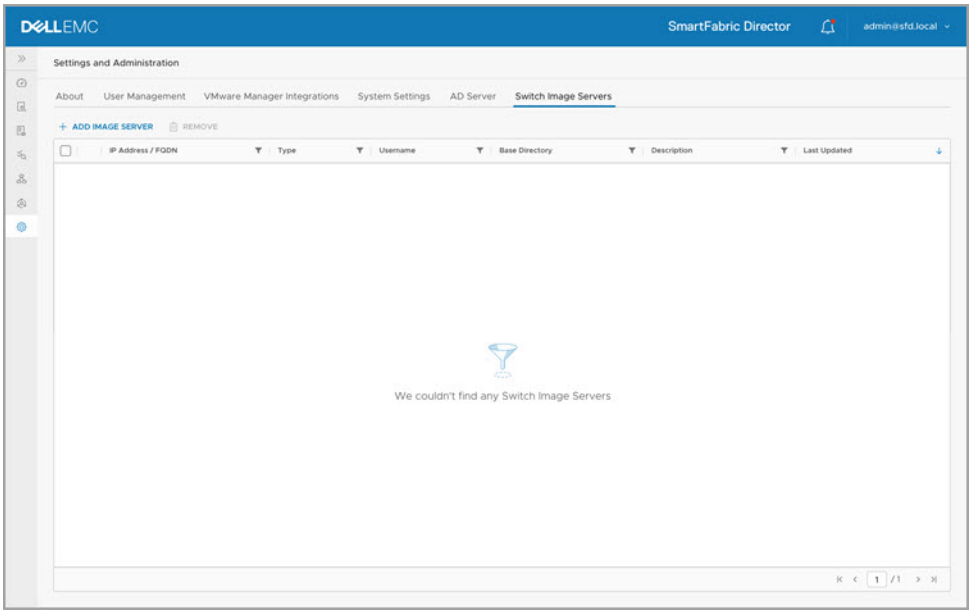

2. Click **Add switch image server**.

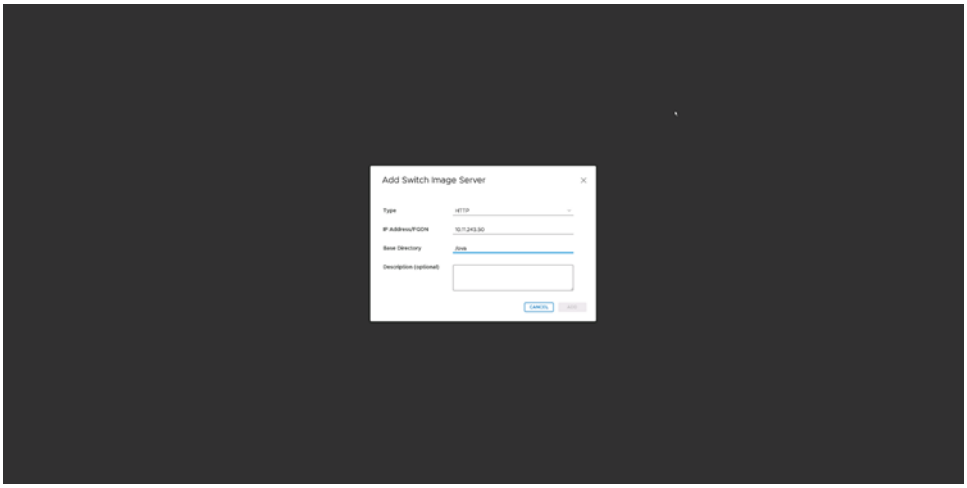

3. Select the image transport type (HTTP, TFTP SFTP, SCP, and FTP), enter the IP address/FQDN, username and password, enter an optional description, then click **Add**.

   The new image server information displays.

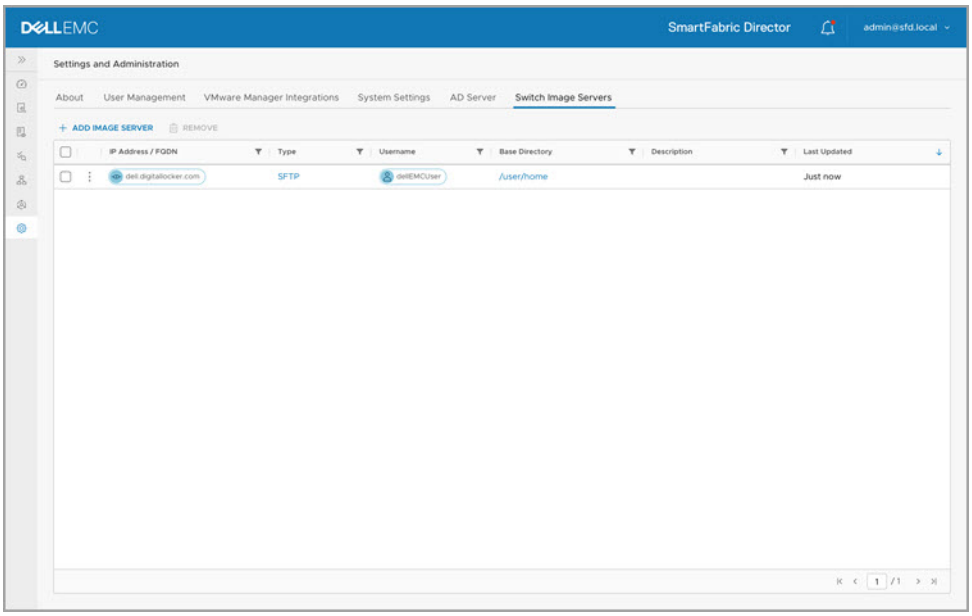

4. (Optional) Select the image server checkbox, then click **Remove** to delete the image server.

Define switch groups

This information describes how to create, edit, and delete switch groups to define an update job. To update a switch image, you must define an update switch lifecycle management job.

You can group switches into a switch group — SFD creates four default switch groups which are autopopulated based on the active fabric wiring diagram.

NOTE: Default switch groups cannot be edited or deleted. These switch groups are automatically created to enable users to upgrade all switches in the predefined switch groups without severely impacting availability of the fabric.

1. Click **New switch group** to define a new switch group.

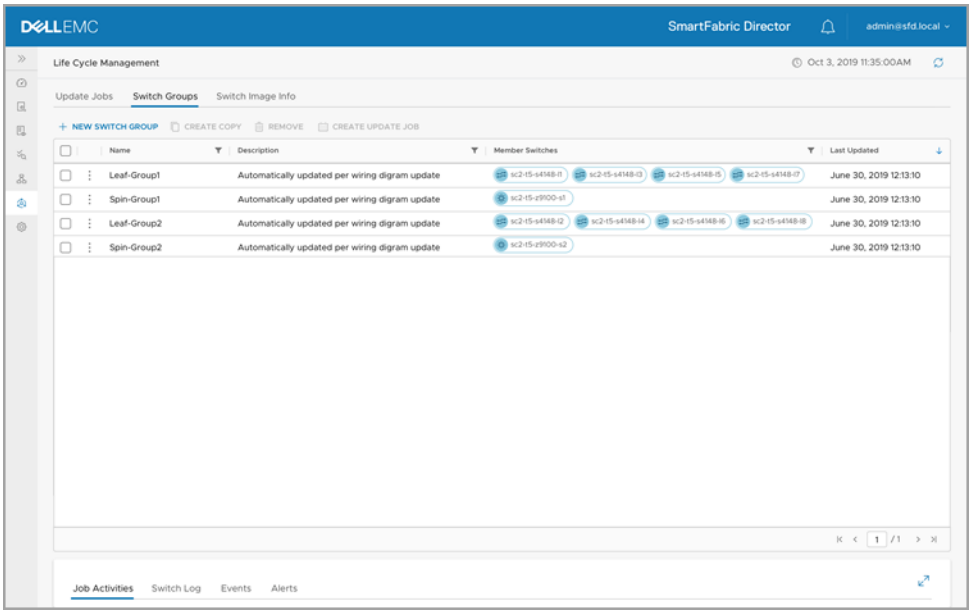

2. Enter the name for the new switch group, select the switches to add to the switch group from the active wiring diagram, enter an optional description, then click **Create**.

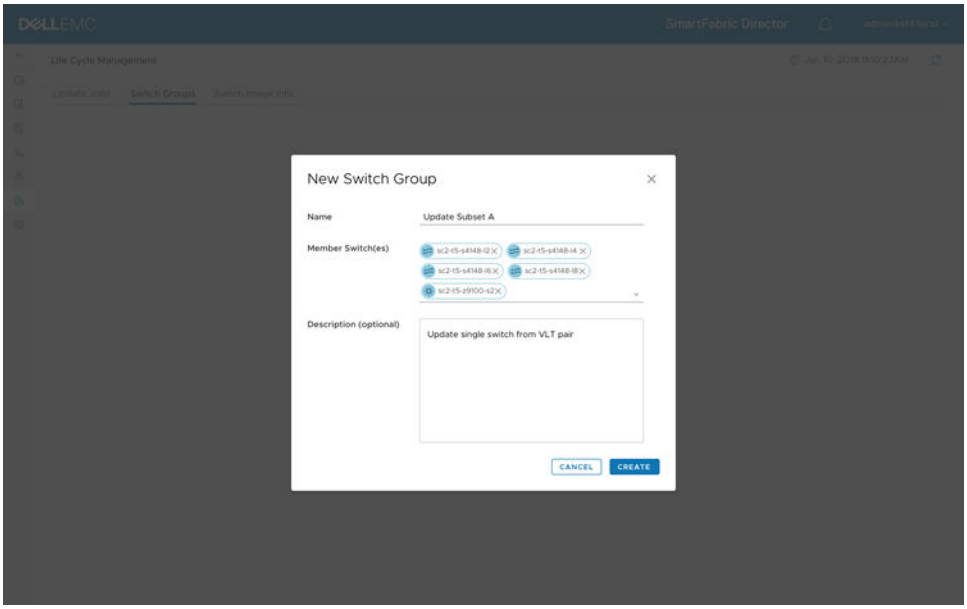

You are now ready to define a switch lifecycle job to create an update job.

Define switch lifecycle job

This information describes how to create an update job. As part of the update job creation, you can select a switch group from the list of available switch groups.

1. Select **Create update job**.

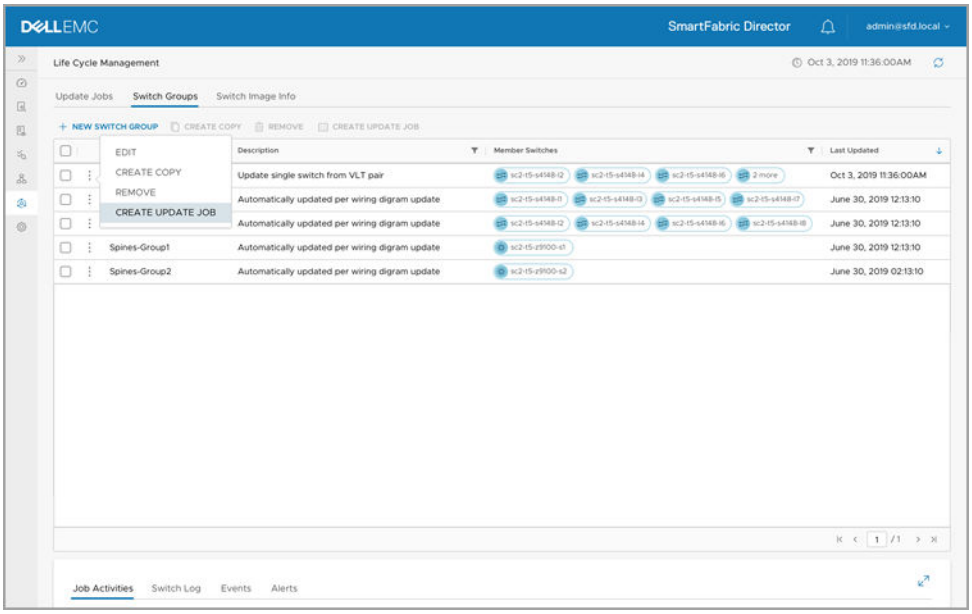

2. Enter the job name, select the switch group, enter an optional description, then click **Next**.

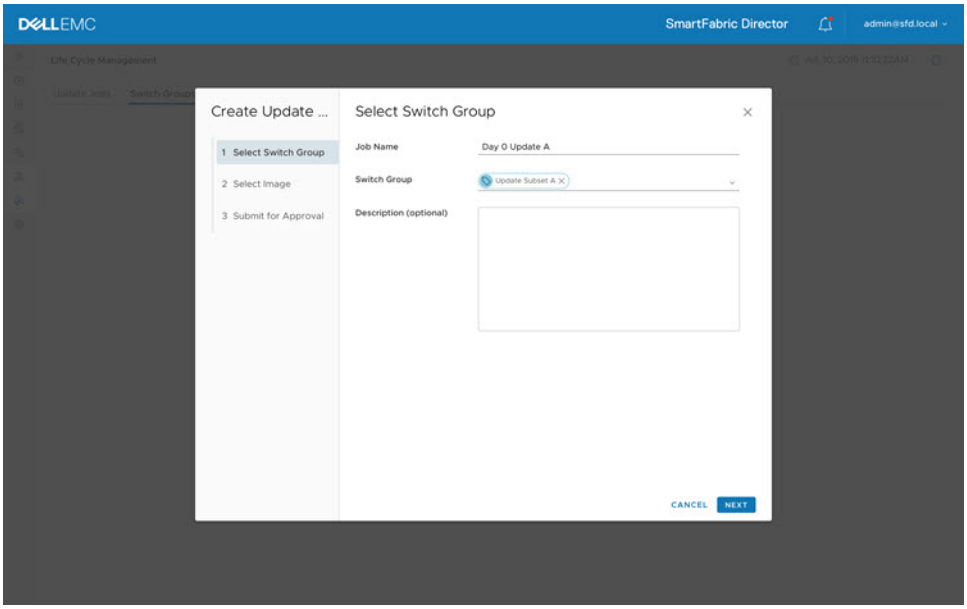

3. Specify an OS10 image including the file extension (.bin). Verify the specified image name matches the name of the image file on the remote server, then click **Next**. All switches in a switch group are updated to the specified image when the Update job is run.

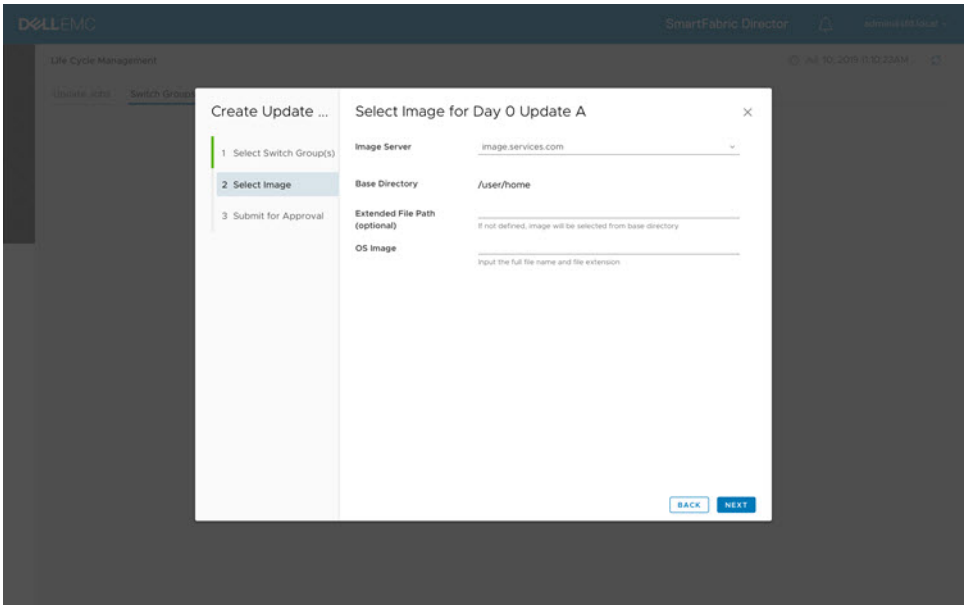

4. Review the image update information, then click **Submit for approval**. You can also click **Save for later** or **Back** to return to the previous screen.

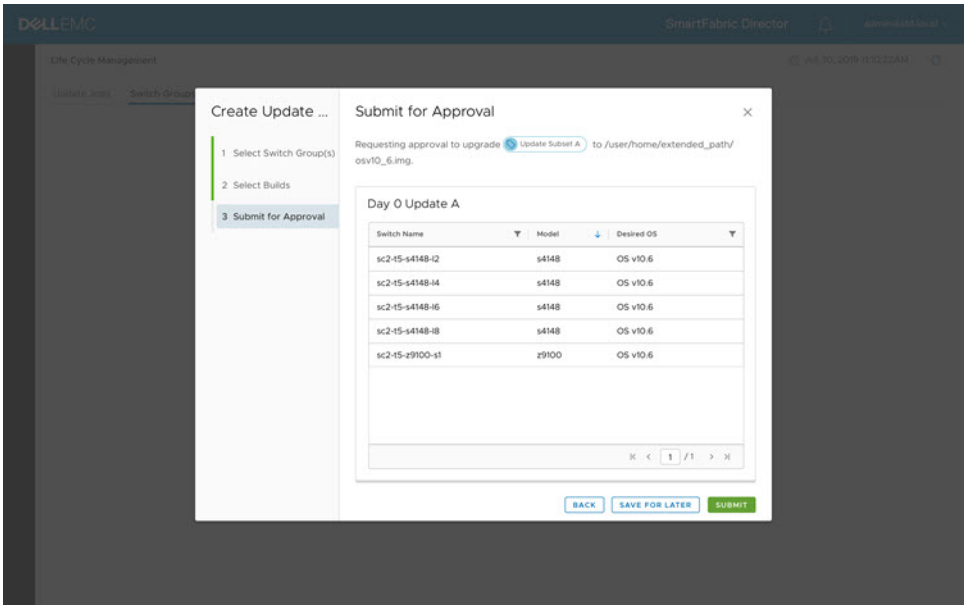

Schedule switch lifecycle job

This information describes how to schedule an Update job. You can schedule an approved job for execution now, or select a future date and time using the calendar.

1. Select **Schedule Now**.

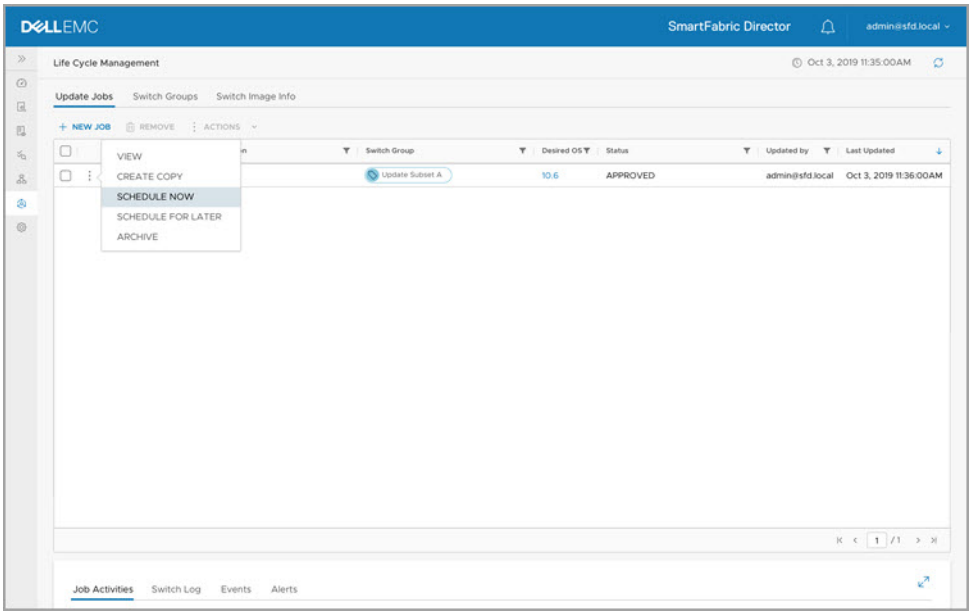

You can also select to schedule the job for a future date and time. Select the date and time, then click **Schedule.**

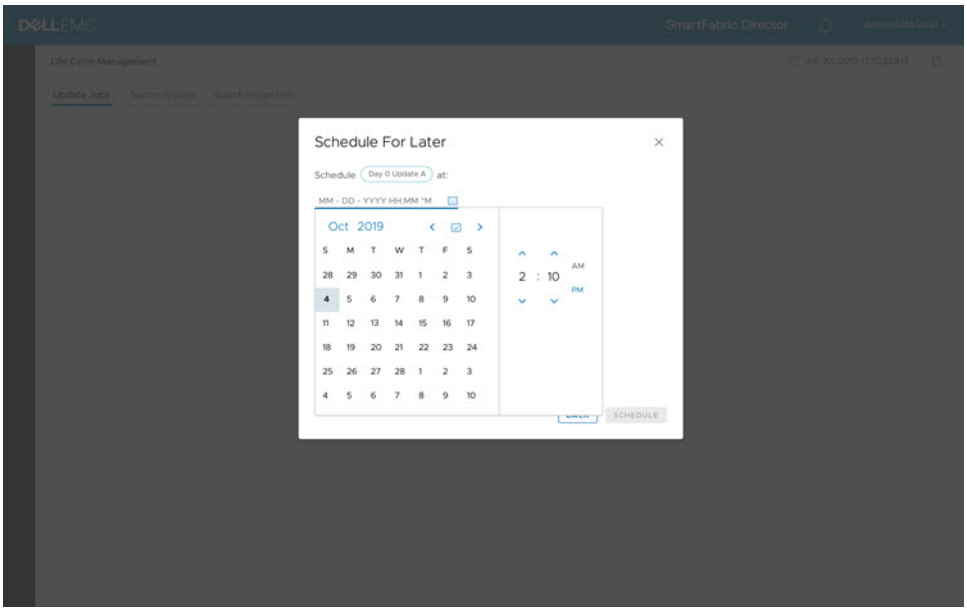

When the job is run, SFD directs the switches to the Image Server to download the specified image. The switch downloads the image, installs the new image, and reboots.

2. Select **View Activity** to display the switch update job activities in the open window.

NOTE: You can also create a copy of the update job to simplify making changes to an existing job.

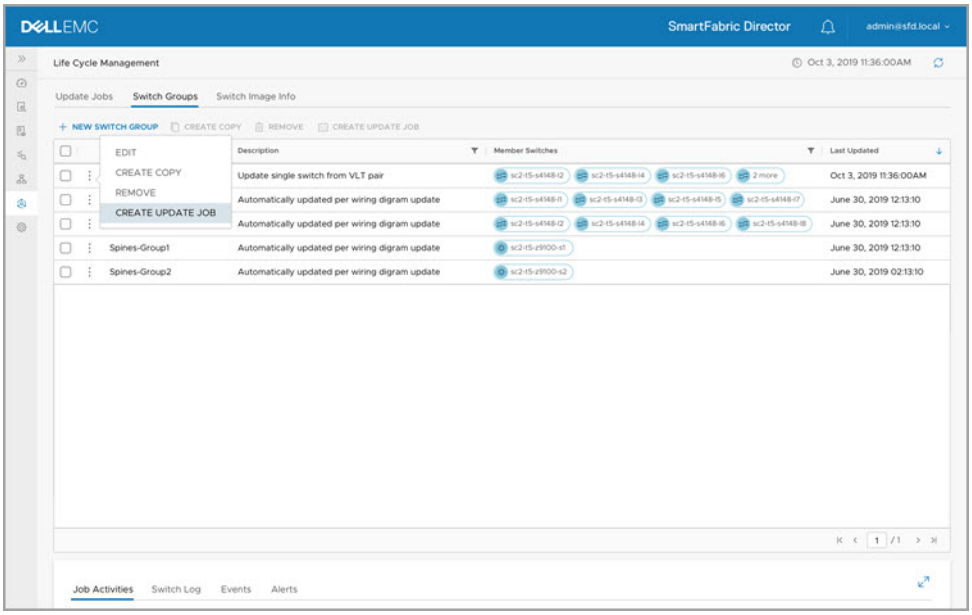

3. Select **View** to display the update job.

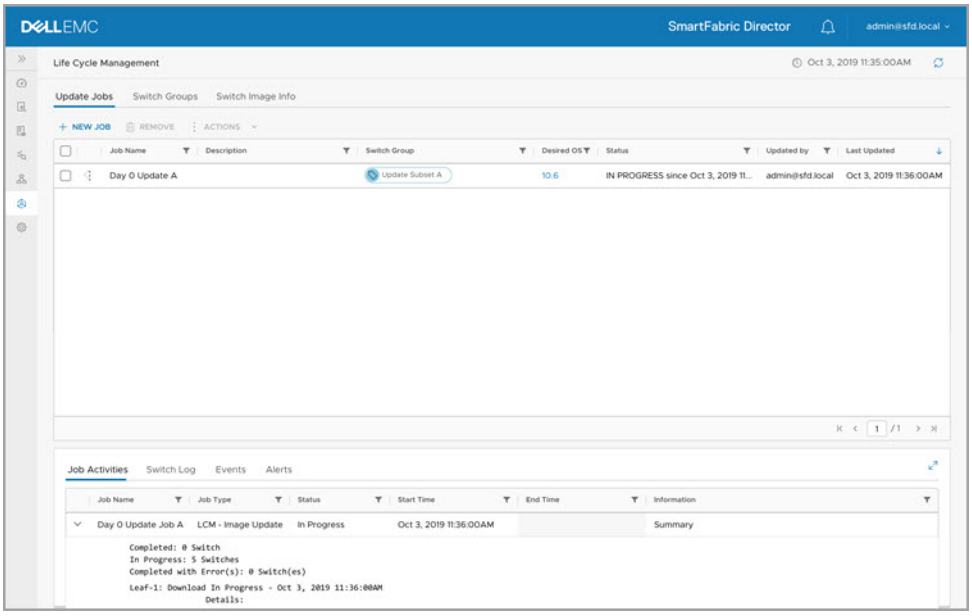

The view displays lifecycle job activities.

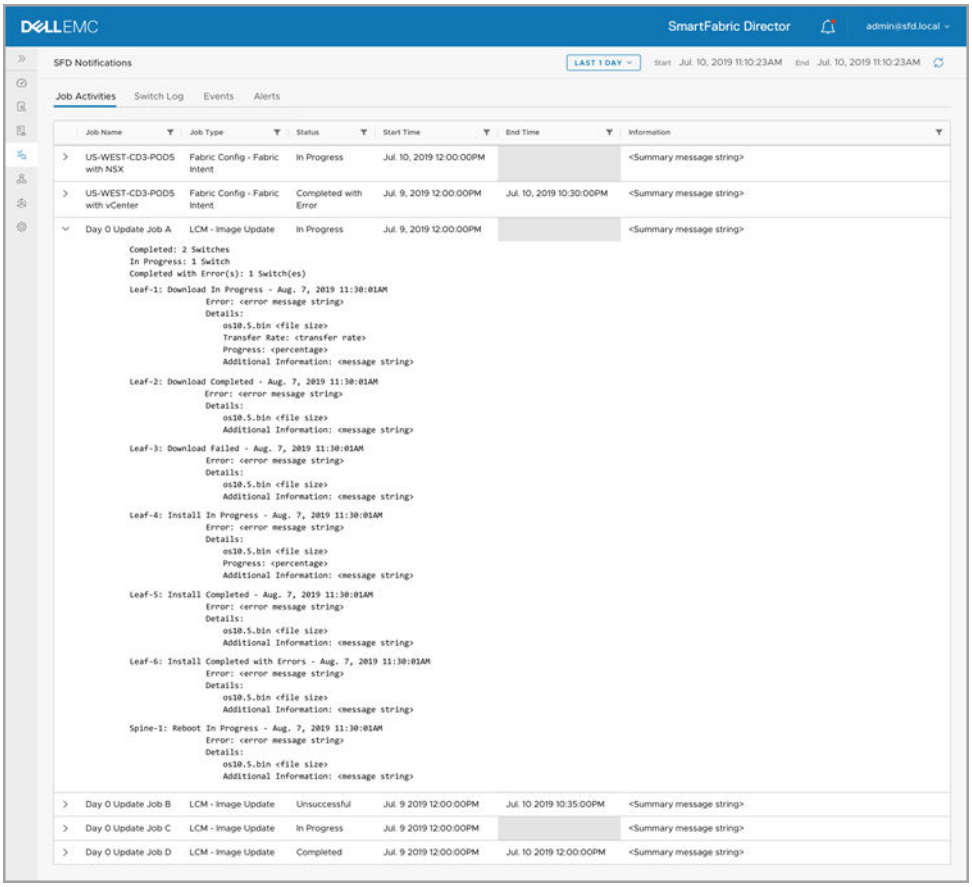
DELL EMC
SmartFabric Director
admin@sfd.local
SFD Notifications
LAST 1 DAY
Start Jul. 10, 2019 11:10:23AM
End Jul. 10, 2019 11:10:23AM
Job Activities
Switch Log
Events
Alerts
Job Name
Job Type
Status
Start Time
End Time
Information
US-WEST-CD3-POD5 with NSX
Fabric Config - Fabric Intent
In Progress
Jul. 10, 2019 12:00:00PM
<Summary message string>
US-WEST-CD3-POD5 with vCenter
Fabric Config - Fabric Intent
Completed with Error
Jul. 9, 2019 12:00:00PM
Jul. 10, 2019 10:30:00PM
<Summary message string>
Day 0 Update Job A
LCM - Image Update
In Progress
Jul. 9, 2019 12:00:00PM
<Summary message string>
Completed: 2 Switches
In Progress: 1 Switch
Completed with Error(s): 1 Switch(es)
Leaf-1: Download In Progress - Aug. 7, 2019 11:30:01AM
Error: <error message string>
Details:
os10.5.bin <file size>
Transfer Rate: <transfer rate>
Progress: <percentage>
Additional Information: <message string>
Leaf-2: Download Completed - Aug. 7, 2019 11:30:01AM
Error: <error message string>
Details:
os10.5.bin <file size>
Additional Information: <message string>
Leaf-3: Download Failed - Aug. 7, 2019 11:30:01AM
Error: <error message string>
Details:
os10.5.bin <file size>
Additional Information: <message string>
Leaf-4: Install In Progress - Aug. 7, 2019 11:30:01AM
Error: <error message string>
Details:
os10.5.bin <file size>
Progress: <percentage>
Additional Information: <message string>
Leaf-5: Install Completed - Aug. 7, 2019 11:30:01AM
Error: <error message string>
Details:
os10.5.bin <file size>
Additional Information: <message string>
Leaf-6: Install Completed with Errors - Aug. 7, 2019 11:30:01AM
Error: <error message string>
Details:
os10.5.bin <file size>
Additional Information: <message string>
Spine-1: Reboot In Progress - Aug. 7, 2019 11:30:01AM
Error: <error message string>
Details:
os10.5.bin <file size>
Additional Information: <message string>
Day 0 Update Job B
LCM - Image Update
Unsuccessful
Jul. 9 2019 12:00:00PM
Jul. 10 2019 10:35:00PM
<Summary message string>
Day 0 Update Job C
LCM - Image Update
In Progress
Jul. 9 2019 12:00:00PM
<Summary message string>
Day 0 Update Job D
LCM - Image Update
Completed
Jul. 9 2019 12:00:00PM
Jul. 10 2019 12:00:00PM
<Summary message string>

10

Using the CLI

This information explains how to access the command-line interface (CLI), and the available commands.

Access the CLI

1. SSH to the IP address configured for SFD.

```
login-srv-05-user%:~> ssh username@sfd.local@ip_address
admin@sfd.local@10.12.124.125's password:
Last login: Mon Oct  17 18:00:59 2019 from 10.12.1.9
```

2. Enter `sfd` to access the SmartFabric Director CLI.

```
admin@sfd.local@SFD-R5:~$ sfd
DellEMC SmartFabric Director CLI
sfd>
```

Command help

To view a list of available options or arguments, enter `-h` or `--help` after any command.

```
sfd> backup --help
usage: backup [-h] {list,create,delete,restore} ...

SFD backup operations - create, delete, list, restore

positional arguments:
  {list,create,delete,restore}

optional arguments:
  -h, --help            show this help message and exit
```

Topics:

- backup
- ftp
- log_level
- service
- support_bundle
- system
- upgrade

backup

Backs up SFD data including the fabric intent.

Command	`backup [list | create --name backup_name | delete --name backup_name | restore --name backup_name]`
Options	• `list` — Displays all current backup files • `create` — Creates a backup

- `delete` — Removes a backup
- `restore` — Restores a backup
- `--name backup_name` — Name of the backup

Usage

Proper backup of SFD is crucial to restore the system to its working state in the event of failure. This feature ensures that all configuration data is backed up. We recommend regular backups — backup frequency and schedule depend on your business needs and operational requirements. At a bar minimum, it is recommended taking backups after any successful deployment, prior to any software upgrades, and after any major Day 2 change. Once you have defined a fabric, configured it and verified that it behaves as expected, use this command to back up all data including the fabric intent. If SFD reaches an undesirable state, you can restore the backup to a golden configuration at any time.

NOTE: The backup and restore versions must match. If you upgrade the software and do not create a backup, you cannot restore an older version backup over a new version. Any configuration changes made between the time a backup was taken, and a restore was performed will be lost. Backup should only be restored on a fresh SFD instance.

1. Log in to the VM to perform restore operations.

```
$ ssh <uname>@<SFD_IP>
```

2. Copy the backup files into the backup directory.

```
$ cp -r <backup_files> /data/sfd_backup/<version_backup>/
```

3. Access the SFD CLI, then view all current backups.

```
$ sfd
sfd> backup list
```

4. Restore the backup file.

```
sfd> backup restore --name sfd_backup
```

5. Close the SFD browser session, then log back into SFD after two to five minutes to start all services.

NOTE: It is recommended to configure FTP (see ftp) to support periodic uploads of backup files.

Examples

sfd> backup list

```
+--------------------------+--------------+--------------------
--+---------------------+----------------+---------------+
| Backup Name | Backup Status | Start Time
| End Time | Start Time (ms) | End Time(ms) |
+--------------------------+--------------+--------------------
--+---------------------+----------------+---------------+
| backup_17-09-2019_08:49:00 | SUCCESS | 17-09-
2019_08:49:00 | 17-09-2019_08:49:01 | 1568710140682 |
1568710141911 |
| test_backup | SUCCESS | 17-09-
2019_14:01:25 | 17-09-2019_14:01:27 | 1568728885977 |
1568728887348 |
| sfd_backup | SUCCESS | 17-09-
2019_14:03:34 | 17-09-2019_14:03:34 | 1568729014086 |
1568729014779 |
+--------------------------+--------------+--------------------
--+---------------------+----------------+---------------+
```

sfd> backup create --name sfd_backup

```
Creating backup sfd_backup
-------------------------------------------------
sfd_backup has been created with status: SUCCESS
-------------------------------------------------
```

Releases

1.1.0 or later

ftp

Configures remote server operations.

Command

`ftp configure --host host_ip --user user --password password --backupdir remote_dir_path --protocol {ftp, sftp} [--port port] [--overwrite {True, False}] ftp list [--frequency {hourly, daily, weekly, monthly} --interval interval]`

Options

- `configure` — Configures the FTP server name and credentials
- `--host host_ip` — Host IP address
- `--user user` — Username
- `--password password` — Password
- `--backupdir remote_dir_path` — Remote directory path to store the file
- `--protocol ftp, sftp` — Either FTP or SFTP for the protocol
- `--port port` — Port number
- `--overwrite True, False` — Either True or False to overwrite the existing file
- `list` — Displays the FTP configuration
- `backup_file` — Name of the backup file to upload
- `--frequency` — Sets the frequency to upload backup files to hourly, daily, weekly, or monthly
- `--interval interval` — Interval value

Usage

This command moves any backup, log, or service pack to an external storage or file server. You must first configure the FTP server name and credentials before uploading files to it from the SFD instance. SFD supports FTP and secure FTP (SFTP) as file transfer protocols. It is recommended that you use periodic upload to optimally use the storage space on the SFD instance. Available frequencies to upload files include hourly, daily, weekly, and monthly.

Example

```
sfd> ftp configure --host 10.196.207.12 --user ftpuser --password
vmware --backupdir public --protocol ftp
Configuring FTP server
Saving FTP parameters
Verifying connection
Creating temp file to verify upload
Connecting to FTP host
Successfully uploaded file to ftp server
Successfully verified connection
Successfully configured ftp server
```

```
sfd> ftp list
+-------------------------------------+
| FTP Configuration                   |
+------------------+------------------+
| Parameter        | Value            |
+------------------+------------------+
| host             | 10.196.207.12    |
| protocol         | sftp             |
| remote_file_path | public           |
| user             | ftpuser          |
+------------------+------------------+
```

```
sfd> ftp upload backup_file --frequency daily
Adding cron job for periodically uploading sfd backups file
Successfully configured periodic upload for backup files
```

Releases

1.1.0 or later

log_level

Sets the log-level for internal events and debug messages.

Command

`log_level list --service service_name all log_level set --service service_name --level log_level`

Options

- `--service service_name` — Lists log-levels for a specific service name
- `all` — Lists log-levels for all available services
- `--level log_level` — Sets the specified service to the wanted log-level (`error`, `warn`, `info`, `debug`, `trace`)

Usage

Use the `all` option to list all service levels.

Example

```
sfd> log_level list --service all
+-----------------------------------------+
| SERVICE LOG LEVELS                      |
+-----------------------------+-----------+
| Service Name                | Log Level |
+-----------------------------+-----------+
| config-builder-service      | INFO      |
| fabric-orchestrator-service | INFO      |
| host-network-service        | INFO      |
| notification-service        | INFO      |
| rest-api                    | INFO      |
| switch-manager-service      | INFO      |
| system-controller-service   | INFO      |
| topology-service            | INFO      |
| telemetry-collector-service | INFO      |
| telemetry-service           | INFO      |
+-----------------------------+-----------+
```

```
sfd> log_level set --service notification-service --level error
+-----------------------------------+
| SET LOG LEVEL OPERATION STATUS    |
+-----------------------+-----------+
| Service Name          | Status    |
+-----------------------+-----------+
| notification-service  | success   |
+-----------------------+-----------+
```

Releases

1.1.0 or later

service

Provides service operations including health and statistics.

Command

`service [list | health --name service_name | stats --name service_name | restart --name service_name]`

Options

- `list` — List all internal services
- `health` — Status of internal services
- `stats` — Metrics including memory and CPU usage of internal services
- `restart` — Restarts the specified service
- `--name service_name` — Service name

Usage

This command provides information about services, performance, and state which can be used for monitoring to diagnose possible problems.

Examples

```
sfd> service list
['config-builder-service', 'elasticsearch',
'external_syslog_collector', 'fabric-orchestrator-service',
'host-network-service', 'infra-processors', 'kube-state-metrics',
'nfc.host', 'nats', 'nats-exporter', 'nats-streaming',
'prometheus', 'prometheus-pgw', 'nginx-exporter', 'nginx-gw',
'node_exporter', 'notification-service', 'rest-api', 'switchmanager-
service', 'system-controller-service', 'topologyservice',
'telemetry-collector-service', 'telemetry-service']
```

```
sfd> service health --name rest-api
-------------------------------------------------
Service-Status : RUNNING
Pod-Status : RUNNING
-------------------------------------------------
```

Releases 1.1.0 or later

support_bundle

Creates a support bundle to be used for debugging purposes.

Command `support_bundle create --name bundle_name`

Options

- `create` — Creates a support bundle
- `--name bundle_name` — Name of the support bundle

Usage This command takes a snapshot of current internal states including health, debug messages, and logging. Verify you have enough local storage before running this command as the file size is large.

NOTE: Do not attempt any SFD operations while the support bundle is being created. Generating a support bundle is CPU intensive, could result in momentary CPU spikes, and may impact the performance of SFD.

Example

```
sfd> support_bundle create --name test
Starting creating support bundle. It will take few minutes to
collect data
Successfully created support bundle test.tar.gz at
/data/nfc_support_bundle/ path
```

Releases 1.1.0 or later

system

Displays the overall software health.

Command `system health`

Options None

Usage None

Example

```
+-------------------------------------------------------+
| SFD-System-Resource                                   |
+-----------------------------------+-----------+-------+
| Name                              | Value     | Unit  |
+-----------------------------------+-----------+-------+
| CPU Usage                         | 14.87     | %     |
| Memory - Available                | 2540158976 | bytes |
| Memory - Available - %            | 30.34     | %     |
```

```
| Available Disk (partition = /)      | 86.44       | Gb    |
| Available Disk (partition = /) - % | 89.34       | %     |
| Network Rate | 146.76 | bps |
+-------------------------------------+-----------+-------+
+-------------------------------------------------------------+
| SFD-System-Health                                           |
+-----------------------------+----------------+------------+
| SFD-Service                 | Service-Status | Pod-Status |
+-----------------------------+----------------+------------+
| config-builder-service      | RUNNING        | RUNNING    |
| elasticsearch               | -              | RUNNING    |
| external_syslog_collector   | RUNNING        | -          |
| fabric-orchestrator-service | RUNNING        | RUNNING    |
| host-network-service        | RUNNING        | RUNNING    |
| infra-processors            | -              | RUNNING    |
| kube-state-metrics          | -              | RUNNING    |
| nfc.host                    | RUNNING        | -          |
| nats                        | -              | RUNNING    |
| nats-exporter               | -              | RUNNING    |
| nats-streaming              | -              | RUNNING    |
| prometheus                  | -              | RUNNING    |
| prometheus-pgw              | -              | RUNNING    |
| nginx-exporter              | -              | RUNNING    |
| nginx-gw                    | -              | RUNNING    |
| node_exporter               | RUNNING        | -          |
| notification-service        | RUNNING        | RUNNING    |
| rest-api                    | RUNNING        | FAILED     |
| switch-manager-service      | RUNNING        | RUNNING    |
| system-controller-service   | RUNNING        | RUNNING    |
| topology-service            | RUNNING        | RUNNING    |
| telemetry-collector-service | RUNNING        | RUNNING    |
| telemetry-service           | RUNNING        | RUNNING    |
+-----------------------------+----------------+------------+
+-----------------------+
| System Overall Health |
+-----------------------+
| DOWN                  |
+-----------------------+
```

Releases 1.1.0 or later

upgrade

Upgrades the SFD software to the specified image.

Command `upgrade` *`bundle_path`* `[--server` *`server_name`*`] [--username` *`username`*`] [--password` *`password`*`]`

Options

- *`bundle_path`* — Path to the bundle host
- `--server` *`server_name`* — IP address of the SCP upgrade bundle host
- `--username` *`username`* — Username of the remote upgrade bundle host
- `--password` *`password`* — SCP password of the remote upgrade bundle host

Usage This command is used to upgrade SFD software by preserving the fabric intent and other configurations across software versions. It is recommended that you create and save a backup of SFD prior to upgrading the software (see backup).

NOTE: The SmartFabric Director software cannot be downgraded to a lower version.

Example

```
sfd> upgrade
```

Releases 1.1.0 or later

11

Frequently asked questions

This information contains answers to frequently asked questions about SmartFabric Director.

Configuration

Do I need to configure the Management interface on each switch?

The Management interface must be configured and enabled on each switch in the fabric (see Management interface).

How do I view switch port profile configuration?

```
OS10# show switch-port-profile 1/1

|  Node/Unit   |  Current    |  Next-boot   | Default    |
|--------------+-------------------+-------------------|
|     1/1      |  profile-2  |  profile-2   | profile-1 |

Supported Profiles:
profile-1
profile-2
profile-3
profile-4
profile-5
profile-6
```

Lifecycle

How do I add a switch group?

See Define switch group for complete information.

Where can I view the status of my image update job?

See Schedule switch lifecycle job for complete information.

Administration

I cannot connect to my image server.

See Specify image servers for complete information.

Maintenance

How can I backup and restore SmartFabric Director?

SmartFabric Director supports back and restore to allow the software to return to a golden configuration at any time. Once the fabric has been defined, configured, and the behavior is verified the operator can use `backup create` to backup the SFD data including the fabric intent. See backup for complete information.

Index

U

V

W