© 2026 BE Networks

Network Hyperautomation

Product Documentation

BE Networks

Copyright BE Networks 2026

Table of contents

1. Introduction 4
1.1 Features 4
2. Architecture 5
2.1 Verity Components 5
2.2 Supported Data Center Topologies 7
3. Quickstart 10
3.1 Quickstart 10
3.2 1. Verity Installation 11
3.3 ONIE and ZTP 39
3.4 Underlay Provisioning 41
3.5 4. Onboarding Devices 42
4. Using Verity 47
4.1 Navigation & UI 47
4.2 Fabrics & Devices 59
4.3 Tenancy 99
4.4 Topology 114
4.5 Templates 117
4.6 Observability 146
4.7 Operations 210
4.8 Administration 253
5. API 289
5.1 Verity API 289
5.2 API Definition 291
6. Reference 292
6.1 How to Use This Documentation 292
6.2 Terminology 293
6.3 Display Icons 297
6.4 Managed Objects 307
6.5 Downloads / Resources 309
6.6 Parameters 311
6.7 END USER LICENSE AGREEMENT 328
7. Changelog 333
7.1 Verity 6.7.0 333
7.2 Verity 6.6.1 334
7.3 Verity 6.6 334

Copyright BE Networks 2026

8. L2 Edge 336

8.1 Verity L2 Edge 336

8.2 Manual Onboarding 337

8.3 Security Lockdown for Switches 341

8.4 Security Lockdown for SDLC/ACS Virtual Machine Applications 343

8.5 Application Reference 347

8.6 Working with ONTs 358

8.7 ZTP and SONiC EDGE Installation 366

9. Getting Help - Report a Problem 377

Copyright BE Networks 2026

1. Introduction

Verity is a software platform for managing physical and logical network infrastructure. The Verity operates data center, campus, edge, and AI network topologies using open-source networking including SONiC and non-proprietary protocols.

1.1 Features

1.1.1 Intent-Based Networking (IBN)

Through a declarative IBN model, Verity enables users to focus on delivering data center services using a simple intent-driven methodology. Verity eliminates the need to convert desired intent into a series of complex, manually entered CLI commands. Instead, the user's intent is captured through a series of intuitive provisioning templates and prompts. Verity automatically maps these inputs into vendor-specific configuration commands that are pushed to the data center fabric with 100% accuracy.

1.1.2 Zero-Touch Provisioning (ZTP)

Verity enables the ZTP installation of SONiC-based switch hardware with only the ONIE boot code installed for modern data center layer 2 and layer 3 fabrics and associated out-of-band switch management networks. Verity allows for the entire site to be pre-designed with no hardware present. A digital twin of the desired network is built within Verity, allowing administrators to design the network underlay according to their unique business requirements. Once the switches are unboxed, racked, powered, and cabled at the data center location, Verity provides the necessary DHCP services so the switches can locate their respective firmware files for automated remote downloading and updating to the newly provisioned hardware. As the ZTP process completes, the switches call into the Verity orchestrator to receive their pre-staged provisioning configurations for automatic underlay creation. The physical network fabric topology is then auto-discovered, reported, and displayed on the GUI site map. An entire data center can be brought online in minutes, not days.

1.1.3 Authoritative Network Source of Truth (NSoT)

Verity acts as an authoritative data repository for all network configuration parameters and physical inventory for data center fabric and out-of-band management networks. It is impossible to achieve high levels of automation and security without this accurate, live view of the intended network state. Representative configuration data collected by the Verity orchestrator includes IP addresses, interface parameters, VLAN mapping, VXLAN configuration, and device locations. Verity inventories all neighboring device details such as hardware make/model, firmware versions, MAC addresses, etc. Additionally, Verity continuously audits the current operational network state against the intended state, with deviations being alerted and automatically remediated.

1.1.4 Live, Zoomable, and Navigable Fabric Map

The Verity UI provides granular network visibility with a live view of the current state of the data center fabric. With just a few clicks and scrolls of a mouse, operators can zoom in and out of networks, pods, and tenants down to the individual switchport level in seconds.

1.1.5 Detailed Real-Time and Historical Actionable Analytics

Up to one year of historical network operational data is maintained in Verity's system automatically. The last 30 days of history is graphically displayed within the Verity UI across varying time periods (30 min., 2 hrs., 12 hrs., 24 hrs. 7 days, and 30 days). Historical statistics and telemetry are available from every attached device, link, port, LAG, optic, switchpoint, and more. All data visible within the UI is easily accessible through Verity's REST API for streamlined third-party monitoring integrations.

1.1.6 Comprehensive Reporting

Verity provides access to a vast number of customizable, filterable, sortable, and exportable reports. All key parameters relating to provisioning, devices, underlay/overlay, ports and more are easy to access and manipulate.

Copyright BE Networks 2026

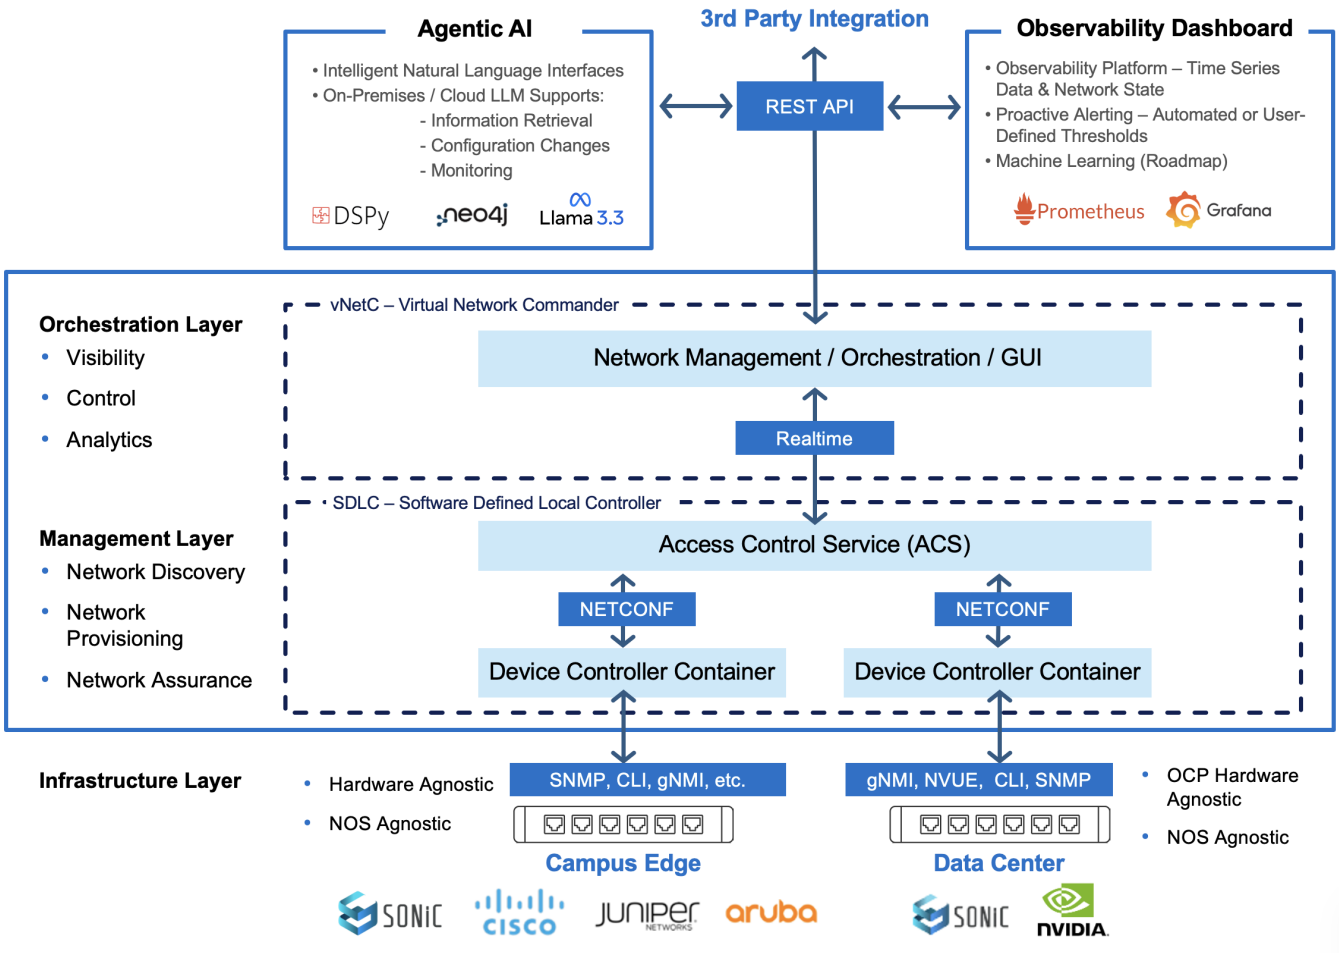
2. Architecture

2.1 Verity Components

2.1.1 Software Defined Local Controller (SDLC)

The SDLC hosts the various containers used for device communications and other system functions. This is a Linux Virtual Machine that is installed during the initial deployment of Verity. This VM hosts the following containers:

- **ACS** (**Access Communication Server**) provides a scalable message interface between vNETC and all managed devices. ACS controls how the devices are programmed from the UI to the Device Controllers.
- **Backend Service(s)** is a system application that increases UI performance. Only one Backend Service can be used per SDLC. This is the UI backend that takes the inputs from the user through the UI and codes it so that the ACS can translate to the Device Controllers and manage the devices.
- **Device Controllers** are used to create the interface between the Verity application and the switch device. Based on the LLDP capabilities of the switch, it can be automatically discovered in the network topology, or it can be statically located.



Copyright BE Networks 2026

2.1.2 Virtual Network Commander

The vNETc hosts database services, a UI HTTPS server, the optional DHCP server, and the BizD application which translates user intent to actual network state. The Web Server also works with the Backend Service(s) on the SDLC's to be able to draw and reflect in real-time what is happening to the devices and configurations.

2.1.3 Monitoring (Observability Suite)

The Verity VM hosts multiple containers to provide network and compute observability. The Monitoring VM gathers telemetry, log, and other metric data for the managed devices in Verity.

2.1.4 Hardware Devices

Verity distinguishes between managed devices and unmanaged devices. Managed Devices are known to the system and mapped to SDLC Device Controllers and can be automatically configured. Unmanaged Devices are devices connected to managed devices but are not configured by Verity, including servers and/or switches.

Supported Devices

2.1.5 Management Network

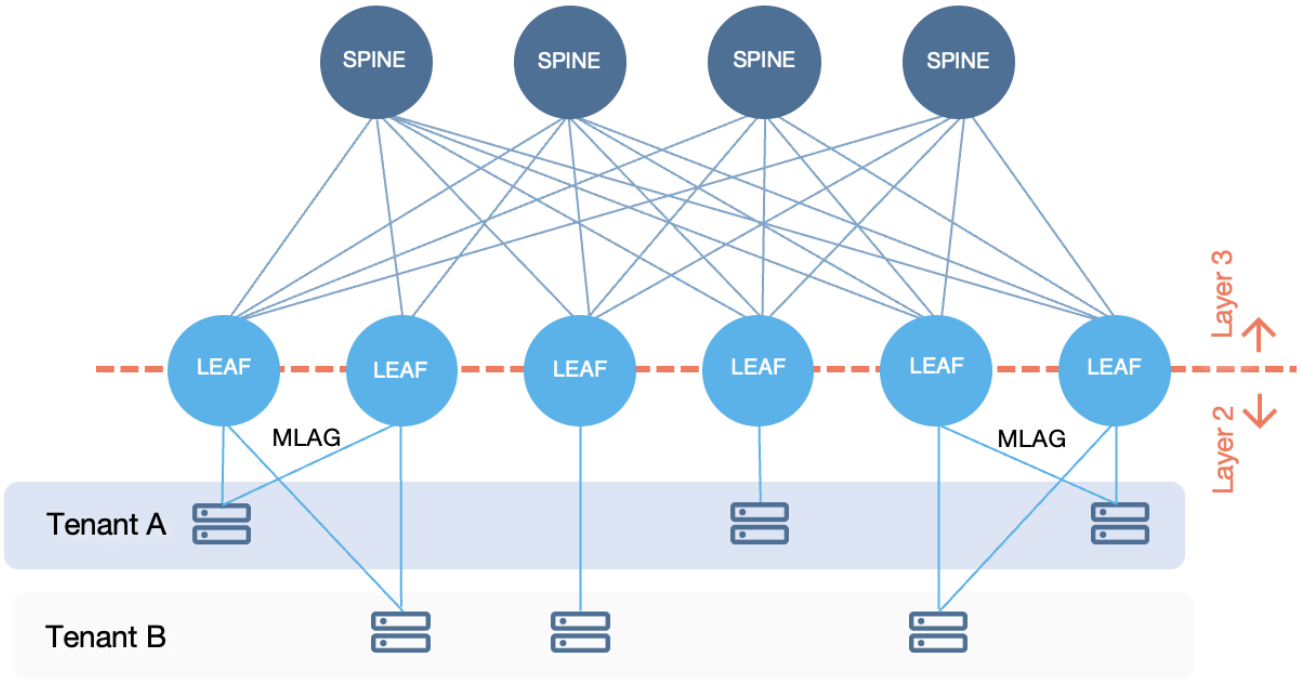
Verity automates the configuration of network devices via the out-of-band ethernet ports connected to a dedicated management network. Verity uses SSH, gNMI, SNMP and other protocols to manage devices in this OOB management network. During initial deployment, Verity can deploy and program the management network and automatically configure the attached switches, as well as inventory and detail how and where they are connected to the management Switch(es).

Copyright BE Networks 2026

2.2 Supported Data Center Topologies

Verity supports a number of different data center topologies, with designs optimized for edge compute all the way up to hyperscale sizes. The core of these designs is a highly redundant load balanced architecture that is widely used in all modern enterprise networks. These designs can be scaled nearly indefinitely, supporting growing network capacity on demand and with low capital expenditures.

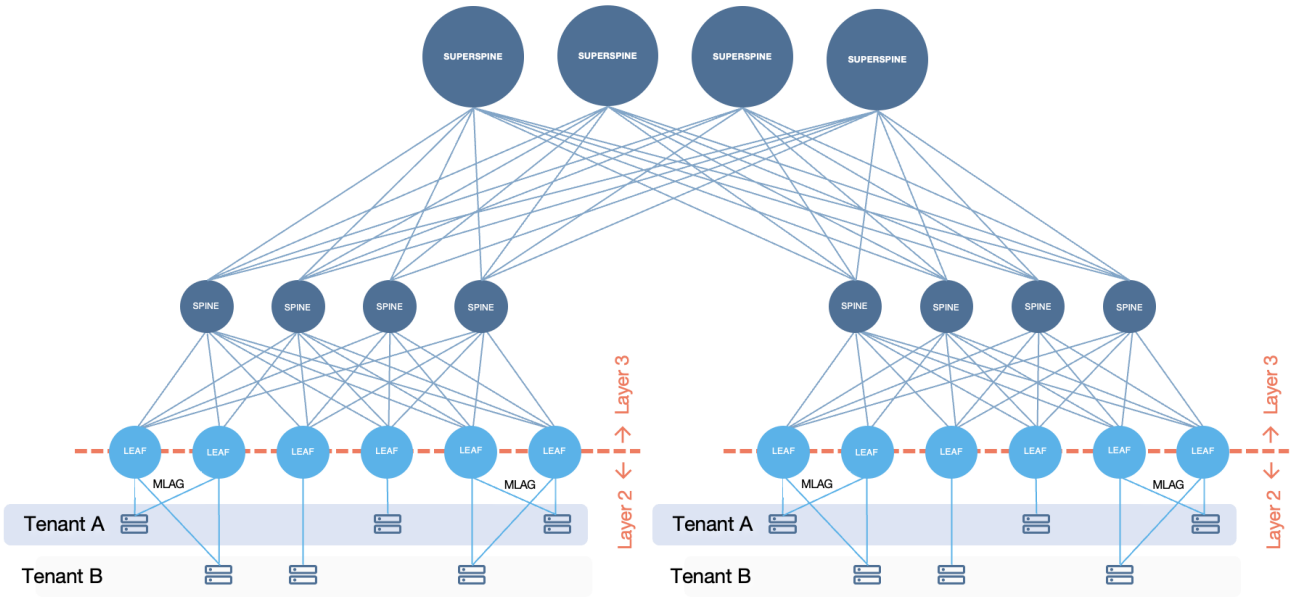
2.2.1 3 Stage Clos Fabric



A 3 Stage Clos Fabric is composed of two layers of switches, generally Leaf and Spine. The leaf switches are typically Top-of-Rack and the spines are either scattered throughout the racks to minimize cable lengths, or positioned at the end of the row (EOR) for simplicity. The 3 stages refers to the path of data from one compute node to another, including the ingress switch, the crossbar switch (spines) and the egress switch. A packet between two compute nodes connected to different leafs will pass through 3 different devices (stages). This topology is commonly referred to as **Leaf-Spine**.

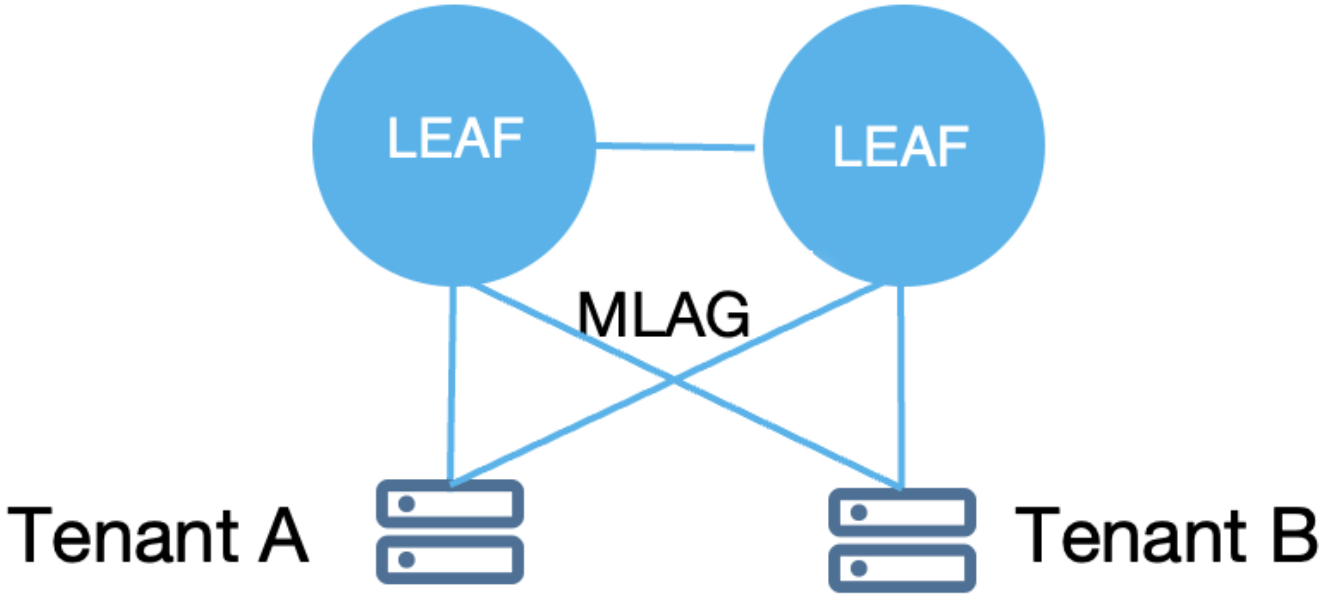
2.2.2 5 Stage Clos Fabric

Copyright BE Networks 2026



A 5 Stage Clos Fabric is identical to a 3 Stage, except that there are multiple groups of leafs and spines. Each group is commonly referred to as a **Pod** and multiple pods are connected through an additional layer of switches known as a Superspine. 5 Stage fabrics typically use BGP route reflectors or route servers to minimize the size of the routing table and simplify device configurations. Only build a 5 stage fabric if you know that your network will be of a size that requires it as the added complexity is typically unnecessary except in ther case of extremely large networks.

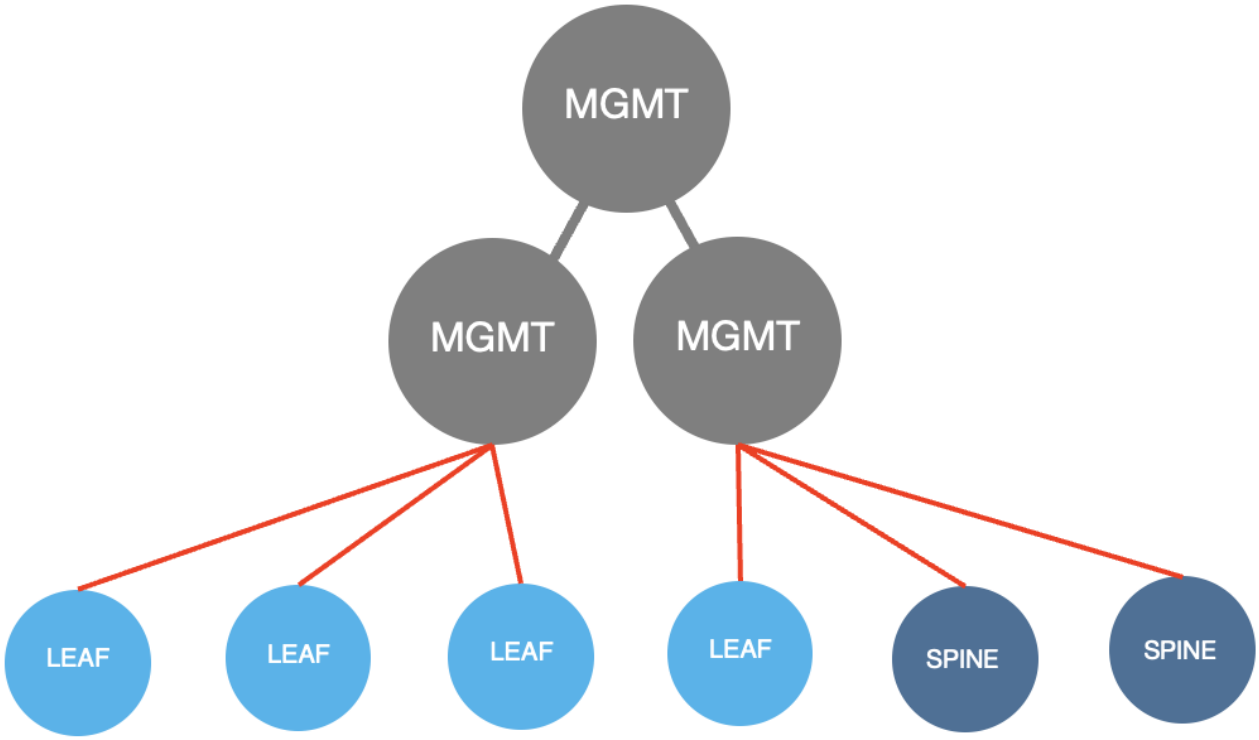
2.2.3 Collapsed Core



A collapsed core design has no spine layers, it is composed of a small number of leaf switches that are all interconnected with one another via MLAG peer connections or a simple Layer-2 trunk between switches. Since there is only one layer, both Layer-2 and Layer-3 functions run on the switches.

2.2.4 Management Network

Copyright BE Networks 2026



Verity supports the management network used to connect to each fabric device for programming and gathering telemetry. Each managed switch is connected typically at 1 Gigabit to a single Layer-2 VLAN extended across a small number of switches.

Copyright BE Networks 2026

3. Quickstart

3.1 Quickstart

Verity EULA Acceptance

Use of this the Verity software requires acceptance of the EULA.

Verity Quickstart has four steps. Upon completing these steps, the administrator will have an operational network fabric with a working underlay network and can begin adding tenants.

1. **Verity Installation** — VMware, KVM, Proxmox, or GNS3 (lab).
2. **Initial Setup**

a. Import an inventory file (FDC*.csv). An *example* FDC inventory import file is available in Downloads.

b. Set up DHCP server and install SONiC images into the platform.

3. **Underlay Settings** — provision the underlay network and reserved ranges.
4. **Onboarding Devices** — power up and connect the SONiC devices. Use bulk FDC import or manual onboarding as needed.

Copyright BE Networks 2026

3.2 1. Verity Installation

3.2.1 Prerequisites

Network Management Architecture

Each system requires a management subnet that can support 5 system IP addresses as well as 3 IP addresses per managed switch. The breakdown is as follows:

IP address Allocations for Management	Component	Allocation
Verity System components	vNETC LAN side, SDLC, ACS, Backend Services, Monitoring	5 Static Addresses
Managed Switches	Verity Switch Controller	1 Dynamic Address per switch
Managed Switches	Switch in ZTP Process	1 Dynamic Address per switch
Managed Switches	Switch post ZTP Process	1 Static Address per switch

1. **vNetC**

a. Resolvable, fully qualified domain name

b. Static IP address, gateway, DNS servers

c. Valid Verity license

2. **SDLC**

a. IP addressing per table above. *This must be routable to the vNetC!*

3. **Controllers (within SDLC)**

a. IP Addressing per table above. The diagram above shows that the controllers are bridged to NIC 2 of the SDLC. *The IP must be on the same VLAN/ subnet as the SDLC!*

4. **Monitoring**

a. IP addressing per table above. *This must be routable to the vNetC!*

5. **Routable or switched network between Verity components and managed switching devices**

a. If using a router or firewall between Verity and the switches, the following ports must be allowed to pass.

i. Port 8080 for gNMI

ii. Port 80 - HTTP (not needed if using TLS option for file downloads)

iii. Port 443 - HTTPS

iv. Port 22 - SSH

v. Port 161 - SNMP

vi. Port 6343 - sFlow

Copyright BE Networks 2026

Hypervisor Configuration (Required)

VMWARE

VMware systems require this additional network configuration:

- **ESXi**

a. Compute resources meeting Verity requirements based on the number of switches being managed. See Resource documentation for computing CPU and memory needs.

b. Virtual Switch 

i. The vNetC and SDLC should be on the same virtual switch in ESXi or at minimum they must be routable.

ii. Your system requires **Promiscuous mode** to be set to **Accept**.

iii. Your system requires **MAC address changes** to be set to **Accept**.

iv. Your system requires **Forged transmits** to be set to **Accept**.

KVM AND PROXMOX

KVM and Proxmox systems require this additional network configuration:

- **KVM**

a. Compute resources meeting Verity requirements based on the number of switches being managed. See Resource documentation for computing CPU and memory needs.

b. Virtual Switch

i. The vNetC, SDLC and Monitoring should be on the same bridge or at minimum they must be routable.

Required Files

Obtain the following files:

Copyright BE Networks 2026

VMWARE REQUIRED FILES

Description	Filename Example	File Type	Notes
vNetC VM Image	vNetC-x_x_x_x.ova	VMware OVA	Resources including vCPU and memory should be adjusted based on iVN resource needs documentation. Networking will need to be altered to the correct virtual switch names used in the server
vNetC "core" Upgrade	core-x_x_x_x- full.tar	tarball	vNetC needs to be updated via GUI SD-Admin immediately after configuration and boot
SDLC VM Image	SDLC-x_x_x_x.ova	VMware OVA	Resources including vCPU and memory should be adjusted based on iVN resource needs documentation. Networking will need to be altered to the correct virtual switch names used in the server
SDLC Binary Firmware Upgrade	firmware-x_x_x_x.tar	tarball	SDLC should be upgraded via web page immediately after configuration and boot
Monitoring VM Image	monitoring_x.x.x.ova	VMware OVA	Resources including vCPU and memory should be adjusted based on iVN resource needs documentation. Networking will need to be altered to the correct virtual switch names used in the server
Monitoring Software Upgrade	monitoring_x.x.x.tar	tarball	Monitoring needs to be updated via GUI SD-Admin immediately after configuration and boot
License	license.cms or sitexxxxx.tar	License file	Is uploaded using GUI
Firmware Upgrade Package	firmware-x_x_x_x.tar	Binary	SDLC should be upgraded via web page immediately after configuration and boot

Copyright BE Networks 2026

KVM & PROXMOX REQUIRED FILES

Description	Filename Example	File Type	Notes
vNetC VM Image	vNetC-x_x_x_x.qcow2	KVM qcow	Resources including vCPU and memory should be adjusted based on resource needs documentation. Networking will need to be altered to the correct bridge names used in the server
vNetC "core" Upgrade	core-x_x_x_x- full.tar	Tarball	vNetC needs to be updated via UI SD-Admin immediately after configuration and boot
SDLC VM Image	SDLC-x_x_x_x.qcow2	KVM qcow	Resources including vCPU and memory should be adjusted based on resource needs documentation. Networking will need to be altered to the correct virtual switch names used in the server
Firmware Upgrade Package	firmware-x_x_x_x.tar	Tarball	SDLC should be upgraded via system upgrader immediately after configuration and boot
Monitoring VM Image	monitoring_x.x.x.qcow2	KVM qcow	Resources including vCPU and memory should be adjusted based on resource needs documentation. Networking will need to be altered to the correct bridge names used in the server
Monitoring Upgrade Package	monitoring_x.x.x.tar	Tarball	Monitoring should be updated via UI AS-Admin immediately after configuration and boot
License	license.cms or sitexxxxx.tar	License file	Uploaded using UI
XML Parameters	vnetc.xml, SDLC.xml and monitoring.xml	xml	Default files are provided and are edited during the installation process. There are 4 available vNETC xml files covering the use cases of single or multiple NICs along with High Availability options

Copyright BE Networks 2026

3.2.2 Verity VM Installation (VMware 7.0.3)

Introduction

The Verity system is comprised of three functional components, all of which are instantiated as Virtual Machines (VMs). This document describes the installation and configuration of these VMs within a VMWare ESXi server.

Prerequisites

- Recommended Network Management Architecture
- VMWare-Specific Network Configuration
- Required Files

Resource Calculator

Use the VM Resource Calculator to determine system resources.

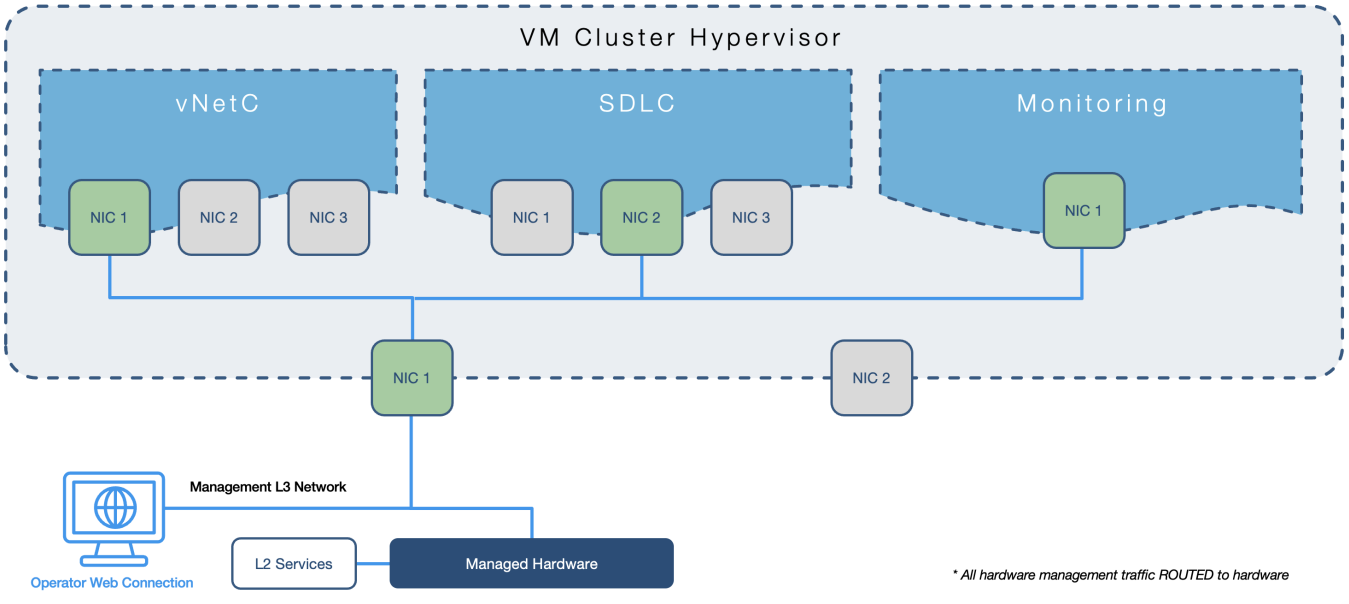
Virtual Machine Overview

Verity's three VM functional components:

Virtual Machine	Function
Virtual Network Commander (vNetC)	Orchestration logic, GUI hosting, northbound RESTful API, and databases.
Software Defined LAN Controller (SDLC)	The SDLC VM is comprised of a series of containers that map one-to-one to the managed switch devices. Network discovery, device provisioning, and network assurance. The SDLC serves as the abstraction layer between the managed switch and the vNetC by translating the native management protocols into the vNetC's NETCONF interface and Yang model.
Monitoring	The Monitoring VM is comprised of various containers that collect, process and display the network device details that are managed by Verity.

VIRTUAL MACHINE TOPOLOGY

Below is the basic VM and hardware topology for reference:

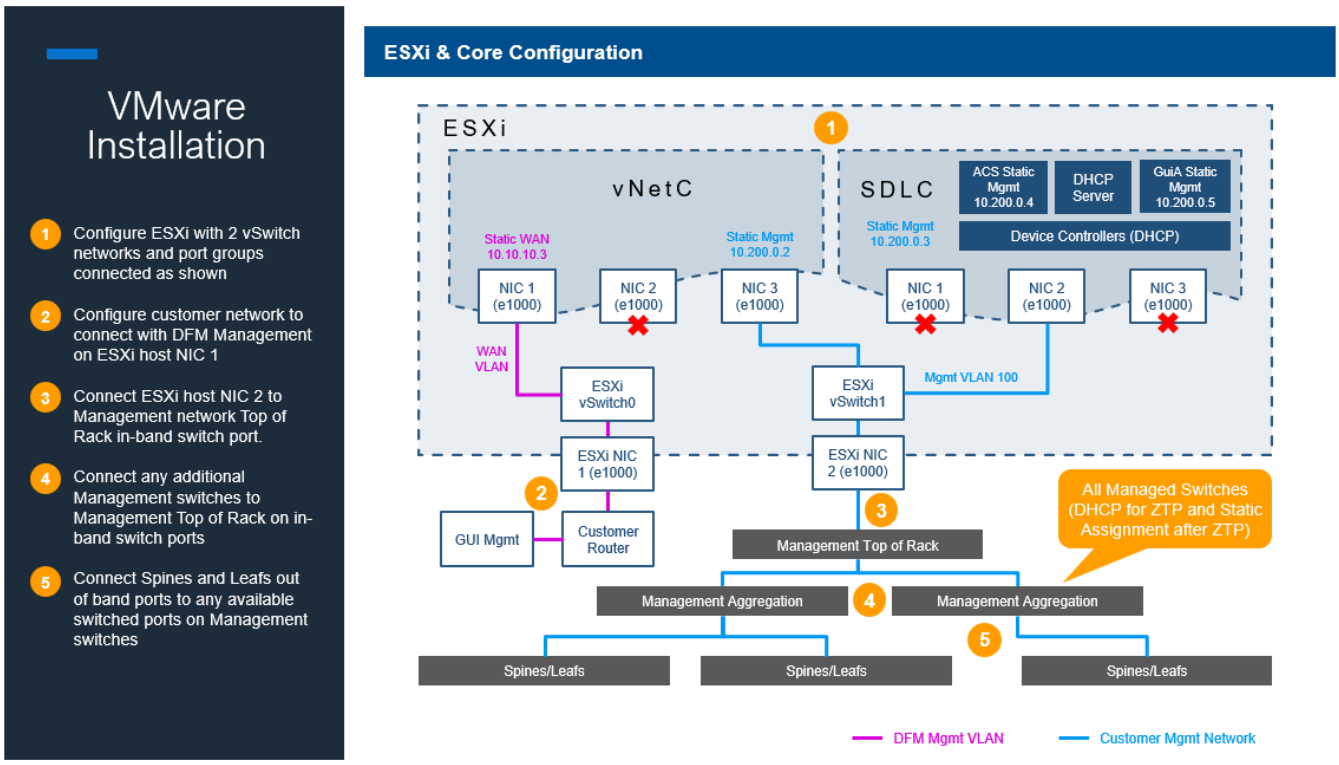


Copyright BE Networks 2026

Public IP Addresses

The standard installation assumes private IP addresses are used by the Verity components and the managed switches. If that is not the case, refer to details in the section "Configure vNetC from the Console" below.

The following diagram shows the recommended management network architecture. Variations are possible based on individual customer's network needs.



Install vNetC VM

1. Go to **Virtual Machines**.
2. Click **Create/Register VM**. 
3. In the window that appears select **Deploy a virtual machine from an OVF or OVA file**. 
4. Click **Next**.
5. Enter a name for the VM and upload the vNetC VM Image OVA file via the prompt that says **Click to Select file or drag/drop**. 
6. Click **Next**.
7. Select the desired data store options. Click **Next**. 
8. Set the Deployment options and network mappings to the correct Port Group.  . This Port Group must be set to promiscuous mode.
9. Click **Next**.
10. Review the settings and if they are correct, click **Finish**. 
11. The VM creation process will start. When the process completes the progress bar in **Recent tasks** at the bottom of the screen will say **Completed Successfully**. 

Copyright BE Networks 2026

CONFIGURE THE VNETC FROM THE CONSOLE

This step requires you to configure the vNetC with an IP address and Fully Qualified Domain Name (FQDN). To do so, you need to open the VM console.

1. Select your VM under the **Virtual Machine** column and click **Console/Open browser console**. 
2. The VM console appears. The vNetc initialization may take several minutes. While waiting you can press **Enter** and wait for login prompt. 
3. Login to the vNetC with username **root** and password **vnc1234**.
4. Enter a new password if prompted (new password = vnc1234). If not prompted for the password, you can continue to use the default password or change it with the **passwd** command.
5. Run the administration application from the shell by typing **ns_admin** and pressing Enter. 
6. You are prompted to enter a web user interface admin account password (ex. admin). Document the password you choose as it will be required for GUI authentication later in the process. It is very important that you remember the password!
7. Press **Enter** when complete.
8. In the Admin Menu, select **Network Configuration**. 
9. Press **Enter**.
10. Select **FQDN (Fully Qualified Domain Name).** Press Enter and set to the desired **Fully Qualified Domain Name.** If the field is prepopulated, it is required that you replace the default text with your own FQDN.
11. Verify that WAN IP DHCP is disabled. If WAN IP DHCP is enabled, disable it using the menu.
12. Select **WAN Static IP Settings,** press enter.
13. Enter the following information:

a. IPv4 Address and subnet in CIDR format (x.x.x.x/#)

b. Default Route (Gateway)

c. DNS Server 1 & DNS Server 2 (optional)

14. Return to the network configuration menu.
15. Save Settings.
16. Follow the prompt and the VM will reboot with the new settings configured.

!!! warning "Public IP Addresses"

```
    The standard installation assumes private IP addresses are use by the Verity components and the managed switches. If that is not the case, refer to the
instructions below.
```

1. ssh back into the vNetC as the root user
2. copy the following string and paste into the command line

ns_vnc_setup --features acs_tunnel=1

Copyright BE Networks 2026

INSTALL THE LICENSE (REQUIRED) AND UPGRADE TO THE LATEST VNETC CORE SOFTWARE

> **Note**
>
> For file transfers to the system for vNETC and other components, users can scp directly to the system as opposed to using file drag-n-drop. Files should be transferred to /var/n/direct location. After successful transfer, the "deploy" button show on the screen as it does after using GUI to transfer files. License file can be directly transferred to the /var/ns/license directory.

1. Use Chrome Web Browser to access the vNetC IP address that was just configured.
2. At the login prompt enter username **admin** and the administration password configured in the menu during installation. These are the credentials you entered in step 3 of **Configure the vNetC from the Console**. 
3. When the window appears, record the information on the **Licensing tied to** line. Provide this information to customer support to obtain your license file. 
4. After you obtain your license.cms file you are required to upload it to the application. In the License window select **data center** or **campus** (depending on your system). Use the drag and drop palette to upload the file or browse for the file. The license file may also be embedded in a \<filename>.tar file and this can also be directly imported and the system will extract the license.cms file. 
5. After you upload the file make sure a success message is presented. 
6. Click the button that says **Yes** to confirm the restart.
7. After the landing page has fully completed populating select the **Admin** button (lower left).
8. Select **Software Packages** tab (top of screen) and click **vNetC Packages**.
9. Using the **Browse Files** (or drag and drop) field, import the **vNetC Core Upgrade** file provided by customer support. 
10. When the process is complete you are presented with a success message. 
11. Click the **Deploy** button.
12. When prompted to continue, click Yes. The software updates. 

> **Temporary Error Message**
>
> You may see an error titled **Fatal Error WebSocket Error: Connection lost -2** appear, this is normal. The browser may temporarily say that the site cannot be reached. When the process is done the landing page will render.

13. If you see a migrations prompt click **Accept**.  .
14. If you see a tan prompt that says **Backend Services not attached, no Backend Service Switch**, clear the message by clicking it.
15. The display should look like the following image: 
16. Go back to the VNC Console in VMware and type `poweroff` in the CLI. This will cleanly shutdown the VNC. 

> **Congratulations**
>
> You have now successfully installed the VNetC VM.

Copyright BE Networks 2026

Install SDLC VM

> **SDLC NIC Mapping when using vCenter**
>
> When deploying the SDLC with vCenter, the NIC to MAC mapping can be remapped, causing connectivity issues with the SDLC and vNetC. If using vCenter to deploy the OVA, check the NIC Mapping on the SDLC by logging in to the SDLC's console, and going to Administration -> Network -> NIC-Map and enter `show` to see mapping:  Verify that the MAC Address for the NIC's in the Mapping match with what the MAC Addresses in VMWare say:  If changes are required, use the `edit` command and select the Index number and map it to the correct MAC Address.

1. Go to **Virtual Machines**. Click **Create/Register VM**. 
2. In the window that appears select **Deploy a virtual machine from an OVF or OVA file**. 
3. Click **Next**.
4. Enter a name for the VM and upload the SDLC VM Image OVA file via the prompt that says **Click to Select file or drag/drop**. Click **Next**. 
5. Select the desired data store options. Click **Next**. 
6. Set the deployment options. Click **Next**. 
7. Review the settings and if they are correct, click **Finish**. 
8. The VM creation process will start. Wait until you see the message **Completed Successfully**. 

Copyright BE Networks 2026

CONFIGURE SDLC (CONSOLE)

The SDLC must be configured with a Static IP address and the vNetC FQDN.

1. Select the SDLC from the VMWARE ESXi interface and click the **Console** tab.
2. Select **Open browser console**. 
3. The console appears. 

> **DHCP Error Messages**
>
> During the following process DHCP errors may appear. These can be ignored.

4. Press **Enter** to get the login prompt, enter username: **admin** and password: **admin**.
5. At the command line interface (CLI) press **Enter** to see a list of options.
6. Select **admin** and press **Enter**.
7. Type **Wizard** and press **Enter**.

> **Note**
>
> If vNetC and SDLC (Backend Services, ACS) are on different subnets, it is recommended to have three consecutive static IP addresses on the same subnet for Backend Services, ACS and DHCP. However, if vNetC is on the same subnet as Backend Services, ACS and DHCP, it is recommended to use four sequential IP addresses.

Prompt	Answer
Enter new hostname	**SDLC**
Enter MGMT IP or enter 'd' to use DHCP	**Enter management IP**
Enter URL connection protocol (http, https)	**http**
Enter default gateway IP/Prefix in CIDR format	**Enter the default gateway IP address**
Enter ACS IP or type 'none' to remove config	**Enter IP**
Enter vNetC FQDN or **IP**	**vNetC IP address**
Enter DNS server	**Enter DNS server IP**
Enter comma separated NTP server(s)	**Enter vNetCs IP address**
Advertise Site Management vlan	**Enter n**
Enter VnetC post SN	**Press Enter**
Enter ACS url	**Press Enter or Enter a different url**

8. Type **y** and press **Enter**
9. Reboot is required for any changes to take effect. In the console, type **reboot** and press **Enter**.

POWER ON THE VNETC

1. In the VMware ESXi interface power on the vNetC. This takes a few minutes.
2. Open the GUI and select select the **Admin Settings** option. 

Copyright BE Networks 2026

3. Set up the **Management VLAN** used to access the Management network. This field is required even if your management switches are using untagged connections.
4. For **Permissible IP Address Ranges on Managed Devices** enter the relevant IP address range (IP address and Mask). 

> **Permissible IP Range Requirement**
>
> The range entered MUST include SDLC components.

5. Click the **Save** button to save your settings. ().
6. Wait until the process is finished. The application landing page resembles the image below when all processes have been completed. 

UPDATE SDLC

> **Note**
>
> For file transfers to the system for SDLC and other components, users can scp directly to the system as opposed to using file drag-n-drop. Files should be transferred to /var/n/direct location. After successful transfer, the "deploy" button show on the screen as it does after using GUI to transfer files.

1. In **Topology/Topology** set the **Upgrades** selection button to the **Enable** setting  .
2. Click the **Admin** tab. Under the **Software Packages** column double-click **Image Packages**.
3. Select and place the SDLC Binary Firmware Upgrade firmware file on the **Drag & Drop** area or use the **Browse Files** button to select the file. 
4. When uploaded, you are prompted with a green success message. 
5. Deploy the upgrade by clicking the **Deploy** button. 
6. A validation message appears. Click **Yes**. 
7. Wait while the package is applied. 
8. Click **Admin** and click **VNFs** .
9. Double click the **SDLC** section. 
10. Double click the box with the title of **SW Version**. 
11. Set the **Target Package** field to the Firmware version  .
12. Click the **Save** button ().
13. Click **Yes** to the validation message.
14. Let the process complete.
15. When the window appears the initial state of **System Applications** are offline. When the **System Applications** come online their LED icons render green. This may take up to 5 minutes. **Admin/VNFs** 

Site Certificate

In order to avoid having to accept the self signed certificate delivered with the system you will need to add a server.pem file to the system. This will need to be obtained from your internet domain administrator.

Copyright BE Networks 2026

The certificate contains the private key, has FQDN set as the common-name, and a certificate authority chain.

1. Go to **Admin**.
2. Under Certificates click **vNetC Server Certificate**. 
3. Drag and drop the **server.pem** file. 

Install Monitoring VM

1. Go to **Virtual Machines**.
2. Click **Create/Register VM**. 
3. In the window that appears select **Deploy a virtual machine from an OVF or OVA file**. 
4. Click **Next**.
5. Enter a name for the VM and upload the vNetC VM Image OVA file via the prompt that says **Click to Select file or drag/drop**. 
6. Click **Next**.
7. Select the desired data store options. Click **Next**. 
8. Click **Next**.
9. Set the Deployment options and network mappings to the correct Port Group.  .
10. This Port Group must be set to **promiscuous mode**.
11. Click **Next**.
12. Review the settings and if they are correct. 
13. Click **Finish**.
14. The VM creation process will start. When the process completes the progress bar in **Recent tasks** at the bottom of the screen will say **Completed Successfully**.

Configure Monitoring VM (Console)

Prerequisites

- **Network Configuration**: You must configure the settings at **Admin/Admin Setting** () to allow access for the subnet where Monitoring will be installed.

Installation & Setup

This step requires you to configure Monitoring with an IP address, default gateway, and DNS servers. Then the script will ask for the Fully Qualified Domain Name (FQDN) of the VNetC so it knows how to connect the Monitoring dashboard. To do so, you need to open the VM console.

1. Select your VM under the **Virtual Machine** column and click **Console**/**Open browser console**. 
2. The VM console appears. The Monitoring initialization may take several minutes. While waiting, you can press **Enter** and wait for login prompt.
3. Login to Monitoring with username **verity** and password **vnc1234**.
4. Enter a new password. Remember this password.
5. You will see the following interface:

```
**************************************
Welcome to the Monitoring Admin Menu
**************************************
```


Copyright BE Networks 2026

```
Please choose an option:
1) setup
2) troubleshooting
#?
```

The two available options are **setup** and **troubleshooting**.

Monitoring Admin Menu

OPTION 1: SETUP

Purpose: Executes the system setup process

When to Use:

- Initial system configuration
- Reinstalling or reconfiguring components
- Setting up new environments

> **Troubleshooting**
>
> For further information about Troubleshooting, visit the Troubleshooting section.

SELECT SETUP

1. When prompted with `#?`, enter the number corresponding to your choice: **1** for `setup`.
2. Press `Enter` to confirm your selection
3. Select management network interface and press Enter.
4. When prompted, enter the following information:

a. Enter static IP Address and subnet in CIDR format (x.x.x.x/#)

b. Default Route (Gateway)

c. DNS Servers separated by a comma

5. Enter the hostname of the server
6. Enter the FQDN or IP address of the vNetC host.
7. Enter TogetherAI key (if applicable).
8. Setup of Monitoring is complete. The display will show the current settings and provide a message on how to make changes in the future.
9. Select y/n to configure email notifications. This is optional.
10. After you set the email notification options you are presented with a prompt that states **Starting Docker Containers**. Wait for this to complete.
11. Setup of Monitoring is complete. The display will show the current settings and provide a message on how to make changes in the future.
12. Type** sudo reboot** to reboot the VM for all the settings to take effect. After the reboot, it takes about 3 minutes for the Docker containers to start up and to announce itself to the vNetC.
13. When the Monitoring VM connects to the vNetC, in Verity, a Growl with the MAC Address of the Monitoring VM will appear. Once this does, use the refresh button on Chrome.
14. There will be a new **Observe** menu option available. Also, the **Observe** Dashboard will be the new startup screen.

    !!! "Upgrade Required"

    ```
    After completing the previous installation steps, the user must install the latest software as outlined below. The installation process is not considered
    complete until this final step is finished.
    ```

15. Use Chrome Web Browser to access the vNetC IP address that was just configured.
16. At the login prompt enter username **admin** and the administration password configured in the menu during installation. These are the credentials you entered in step 3 of **Configure the vNetC from the Console**.

Copyright BE Networks 2026

17. From the **Admin** web page and under the **Software Packages** heading, click **App Packages**. 
18. Using the **Browse Files** (or drag and drop) field, import the **Monitoring Upgrade** file provided by customer support. 
19. When the process is complete you are presented with a success message. 
20. Click the **Deploy** button.
21. When prompted to continue, click Yes. The software updates. 

> **Wait for System**
>
> It take about 5 minutes for the tarball to be uploaded to the Monitoring VM, and the changes to be applied and the new containers started up. If you SSH into the Monitoring VM, and go to the /be_install directory, you will see the tarball uploaded. If you run sudo docker ps you will see the uptime of the containers to be less than 5 minutes online letting you know that everything is updated.

> **Congratulations**
>
> Verity has been successfully installed. Treat yourself to an iced coffee!

Copyright BE Networks 2026

3.2.3 KVM Installation

Introduction

The Verity system is comprised of three functional components, all of which are instantiated as Virtual Machines (VMs). This document describes the installation and configuration of these VMs within a KVM environment.

Prerequisites

- Recommended Network Management Architecture
- KVM and Proxmox Specific Network Configuration
- KVM & Proxmox Required Files

Resource Calculator

Use the VM Resource Calculator to determine system resources.

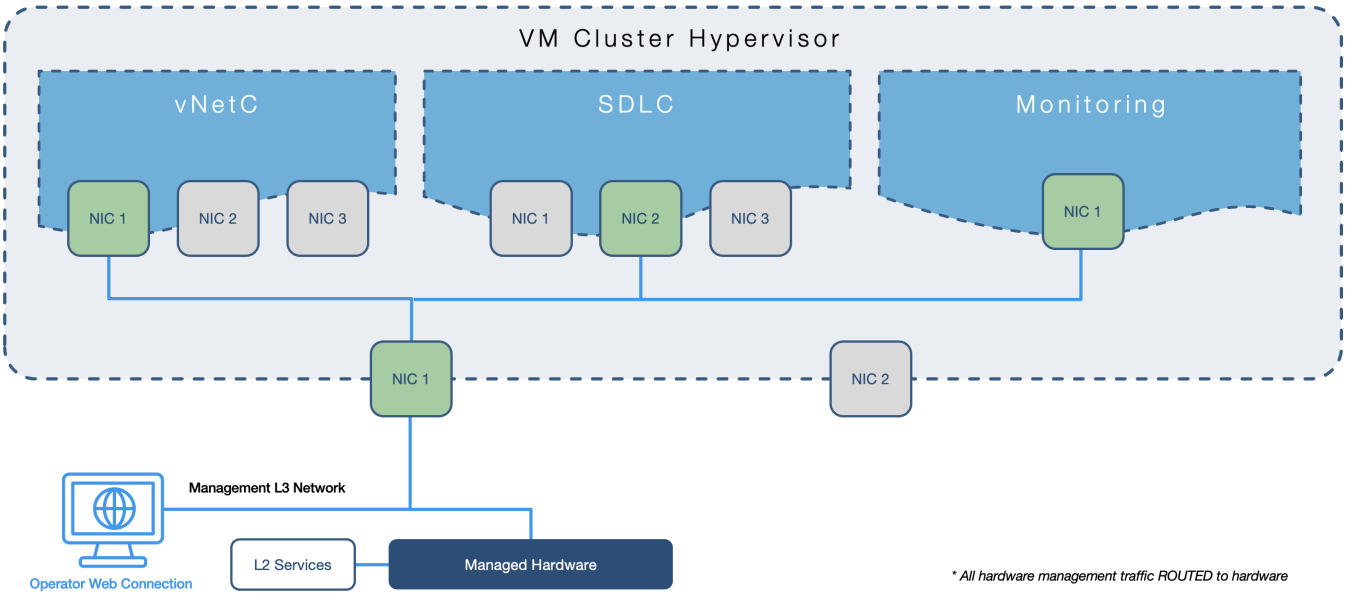
Virtual Machine Overview

Verity's three VM functional components:

Virtual Machine	Function
Virtual Network Commander (vNetC)	Orchestration logic, GUI hosting, northbound RESTful API, and databases.
Software Defined LAN Controller (SDLC)	The SDLC VM is comprised of a series of containers that map one-to-one to the managed switch devices. Network discovery, device provisioning, and network assurance. The SDLC serves as the abstraction layer between the managed switch and the vNetC by translating the native management protocols into the vNetC's NETCONF interface and Yang model.
Monitoring	The Monitoring VM is comprised of various containers that collect, process and display the network device details that are managed by Verity.

Topology Overview

Below is the basic VM and hardware topology for reference:



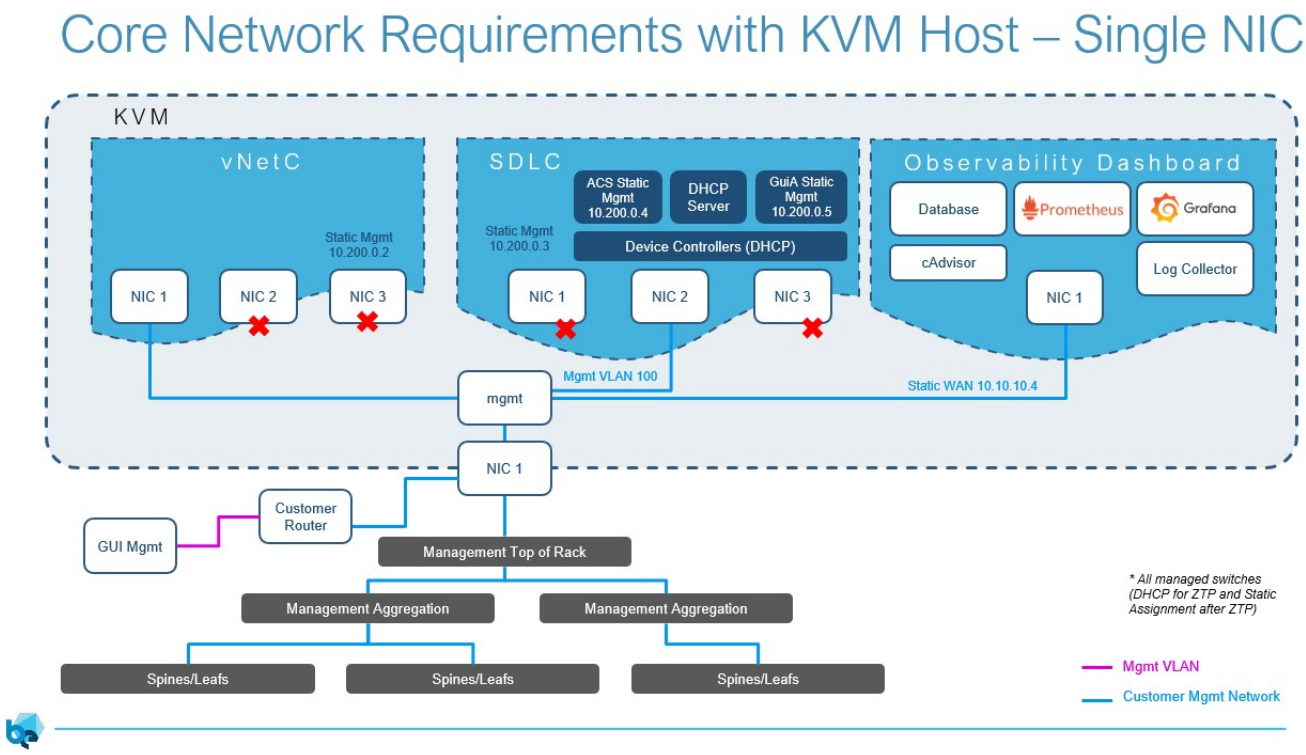
Copyright BE Networks 2026

The orchestration platform (vNETC) is configured on the customer's network with one static IP address to be accessed by users.

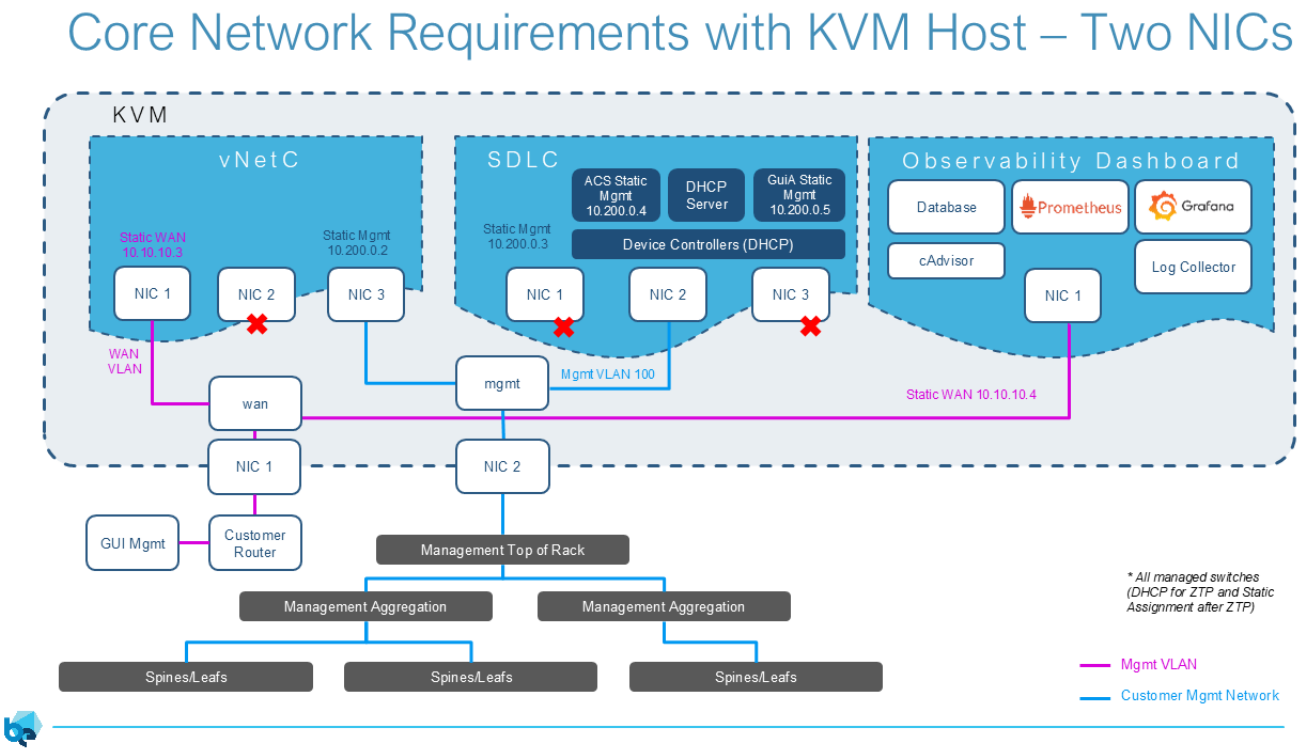
Public IP Addresses

The standard installation assumes private IP addresses are use by the Verity components and the managed switches. If that is not the case, refer to details in the section "Configure vNetC from the Console" below.

The following diagrams shows the recommended management network architectures. Variations are possible based on individual customer's network needs. The VMs can have all of their interfaces on a single NIC coming from the server, or a secondary NIC can be used to separate the WAN connection to the vNETC from the Management network connecting to all of the switches.



Copyright BE Networks 2026



3.2.4 Creating the Virtual Machines

The following instructions explain how to create virtual machines for both the vNetC and SDLC.

1. Copy vnetc, sdlc and monitoring qcow and xml files to host (root directory):

a. vnetc.qcow2

b. SDLC.qcow2

c. monitoring.qcow2

d. vnetc.xml

e. SDLC.xml

f. monitoring.xml

2. Make sure you have the bridge name from host for VMs: `

```
nmcli connection show
```

`

3. Edit xml files (vnetc.xml, SDLC.xml and monitoring.xml) to have bridge name on the correct interfaces, number of CPU, and number RAM needed for each VM. Example:

```
<memory unit="KiB">8388608</memory>
<currentMemory unit="KiB">8388608</currentMemory>
<vcpu placement="static">8</vcpu>

<interface type="bridge">
<source bridge="mgmt"/>

<interface type="bridge">
```


Copyright BE Networks 2026

```
<source bridge="wan"/>
<model type="virtio"/>
```

1. Adjust xml file for Linux variants

a. For Centos and Redhat

`

```
<devices>
<emulator>/usr/libexec/qemu-kvm</emulator>

<os>
<type arch='x86_64' machine='pc-i440fx-rhel7.6.0'>hvm</type>
```

`

a. For Ubuntu!!! note If vNetC and SDLC (Backend Service(s), ACS) are on different subnets, it is recommended to have three consecutive static IP addresses on the same subnet for Backend Service(s), ACS and DHCP.

`

```
<devices>
<emulator>/usr/bin/qemu-system-x86_64</emulator>

<os>
<type arch='x86_64' machine='pc-i440fx-jammy'>hvm</type>
```

`

2. Move qcow (vnetc, SDLC and monitoring) to /var/lib/libvirt/images directory
3. Define VMs using the edited xml files:

   virsh define vnetc.xml

   (Domain 'vnetc' defined from vnetc.xml)

   virsh define SDLC.xml

   (Domain 'SDLC' defined from SDLC.xml)

   virsh define monitoring.xml

   (Domain 'monitoring' defined from monitoring.xml)
4. Set VMs to autostart on boot:

   virsh autostart vnetc

   (Domain 'vnetc' marked as autostarted)

   virsh autostart SDLC

   (Domain 'SDLC' marked as autostarted)

   virsh autostart monitoring

   (Domain 'monitoring' marked as autostarted)
5. Start vnetc VM

   virsh start vnetc

   (Domain 'vnetc' started)

3.2.5 Configure the vNetC (Console)

This step requires you to configure the vNetC with an IP address and Fully Qualified Domain Name (FQDN). To do so, you need to open the VM console: **virsh console vnetc**

Copyright BE Networks 2026

The VM console appears. The vNetc initialization may take several minutes. While waiting you can press **Enter** and wait for login prompt.

1. Login to the vNetC with username **admin** and password **vnc1234**. Enter a new password if prompted. If not prompted for the password, you can continue to use the default password or change it with the **passwd** command.
2. In the Admin Menu, select **Network Configuration**. Press Enter. 
3. Select **FQDN (Fully Qualified Domain Name).** Press Enter and set to the desired **Fully Qualified Domain Name.** If the field is prepopulated, it is required that you replace the default text with your own FQDN.
4. Verify that WAN IP DHCP is disabled. If WAN IP DHCP is enabled, disable it using the menu.
5. Select **WAN Static IP Settings,** press enter.

- Enter: IPv4 Address and subnet in CIDR format (x.x.x.x/##) where x.x.x.x is the IPv4 address and ## is the CIDR subnet mask prefix
- Enter: Default Route (Gateway)
- Enter: DNS Server 1
- Enter: DNS Server 2 (if required)

6. Return to the network configuration menu.
7. Save Settings

   Follow the prompt and the VM will reboot with the new settings configured.

> **Public IP Addresses**
>
> The standard installation assumes private IP addresses are use by the Verity components and the managed switches. If that is not the case, refer to the instructions below.

1. ssh back into the vNetC as the root user
2. copy the following string and paste into the command line

   ns_vnc_setup --features acs_tunnel=1

Copyright BE Networks 2026

INSTALL THE LICENSE (REQUIRED) AND UPGRADE TO THE LATEST VNETC CORE SOFTWARE

> **Note**
>
> For file transfers to the system for vNETC and other components, users can scp directly to the system as opposed to using file drag-n-drop. Files should be transferred to /var/n/direct location. After successful transfer, the "deploy" button show on the screen as it does after using GUI to transfer files.

1. Use Chrome Web Browser to access the vNetC IP address that was just configured.
2. At the login prompt enter username **admin** and the administration password configured in the menu during installation. These are the credentials you entered in step 3 of **Configure the vNetC from the Console**. 
3. When the window appears, record the information on the **Licensing tied to** line. Provide this information to customer service to obtain your license file. 
4. After you obtain your license.cms file you are required to upload it to the application. In the License window select **data center** or **campus** (depending on your system). Use the drag and drop palette to upload the file or browse for the file. The license file may also be embedded in a \<filename>.tar file and this can also be directly imported and the system will extract the license.cms file. 
5. After you upload the file make sure a success message is presented. 
6. Click the button that says **Complete**.
7. After the **Verity** window has fully completed populating select the **Admin** tab.
8. From the **Admin** option, under **Software Packages** click **vNetC Packages**.
9. Using the **Browse Files** (or drag and drop) field, import the **vNetC Core Upgrade** file. 
10. When the process is complete you are presented with a success message. 
11. Click the **Deploy** button.
12. When prompted to continue, click Yes. The software updates. 

> **Temporary Error Message**
>
> You may see an error titled **Fatal Error WebSocket Error: Connection lost -2** appear, this is normal. The browser may temporarily say that the site cannot be reached. When the process is done the landing page will render.

1. If you see a migrations prompt click **Accept**.  .
2. If you see a tan prompt that says **Backend Service(s) not attached, no Backend Service Switch**, clear the message by clicking it.
3. The display should look like the following image: 
4. Press `<ctrl> ]` in the vNetc console window to return to the host server. Then, stop the vNetc VM by entering `virsh shutdown vnetc`.

3.2.6 Configure the SDLC (Console)

The SDLC must be configured with a Static IP address and the vNetC FQDN. Start the SDLC VM by entering `virsh start SDLC` in the host console. Then, log in to the SDLC console by entering `virsh console SDLC`.

Copyright BE Networks 2026

> **DHCP Error Messages**
>
> During the following process DHCP errors may appear. These can be ignored.

1. Press **Enter** to get the login prompt, enter username: **admin** and password: **admin**.
2. At the command line interface (CLI) press **Enter** to see a list of options.
3. Select **Admin** and press **Enter**.
4. Type **Wizard** and press **Enter**.

> **Note**
>
> If vNetC and SDLC (Backend Service(s), ACS) are on different subnets, it is recommended to have three consecutive static IP addresses on the same subnet for Backend Service(s), ACS and DHCP. However, if vNetC is on the same subnet as Backend Service(s), ACS and DHCP, it is recommended to use four sequential IP addresses.

Prompt	Answer
Enter new hostname	**SDLC**
Enter MGMT IP or enter 'd' to use DHCP	**Enter management IP**
Enter URL connection protocol (http, https)	**http**
Enter default gateway IP/Prefix in CIDR format	**Enter the default gateway IP address**
Enter ACS IP or type 'none' to remove config	**Enter IP**
Enter vNetC FQDN or **IP**	**vNetC IP address**
Enter DNS server	**Enter DNS server IP**
Enter Comma separated NTP server(s)	**Enter vNetCs IP address**
Advertise Site Management vlan	**Enter n**
Enter VnetC post SN	**Enter N/A**
Enter ACS url	**Press Enter or Enter a different url**

1. Type **y** and press **Enter**
2. Reboot is required for any changes to take effect. In the console, type **reboot** and press **Enter**.
3. Exit the SLDC VM using `<ctrl> ]`

3.2.7 Configure Management Network

1. Power on the vNetc VM using this command: `virsh start vnetc`.
2. Open the GUI and select **Admin** in the lower left.
3. Under **Network** select **Admin Settings**. 
4. Set up the **Management VLAN** used to access the Management network. This field is **required** even if your management switches are untagged connections. **Note: If untagged select “Managed on Native VLAN” checkbox**.
5. In **Permissible IP Address Ranges on Managed Devices** enter the relevant IP address range (IP address and Mask). 
6. Click the checkbox icon () to save your settings.

Copyright BE Networks 2026

7. Go to the Topology navigation window. If a beige notification box appears in the lower right, click it to close it.
8. Wait until the process is finished. The application landing page resembles the image below when all processes have been completed. 

3.2.8 Update SDLC

> **Note**
>
> For file transfers to the system for SDLC and other components, users can scp directly to the system as opposed to using file drag-n-drop. Files should be transferred to /var/n/direct location. After successful transfer, the "deploy" button show on the screen as it does after using GUI to transfer files.

1. In **Topology** from **Network View** uncheck the box titled **Disable Upgrades**. 
2. Click the **Admin** tab. Under **Software Packages** click **Image Packages**.
3. Select and place the SDLC Binary Firmware Upgrade firmware file on the **Drag & Drop** area or use the **Browse Files** button to select the file. 
4. When uploaded, you are prompted with a green success message. 
5. Deploy the upgrade by clicking the **Deploy** button. 
6. A validation message appears. Click **Yes**. 
7. Wait while the package is applied. 
8. Click **Admin** and then click **VNFs**. 
9. Double click the **SDLC** section. 
10. Double click the box with the title of **SW Version**. 
11. Set the **Target Package** field to the Firmware version. 
12. Click the **Save** button. ()
13. Click **Yes** to the validation message. 
14. Let the process complete.
15. When the window appears the initial state of **System Applications** are offline. When the **System Applications** come online their LED icons render green. This may take up to 5 minutes. 

3.2.9 Site Certificate

In order to avoid having to accept the self signed certificate delivered with the system you will need to add a server.pem file to the system. This will need to be obtained from your internet domain administrator.

1. Go to **Admin** and under **Certificates** click on **vNetC Server Certificate** box. 
2. Drag and drop the **server.pem** file. 

The installation and updates of Verity components is now complete.

3.2.10 Create the Monitoring VM

Copyright BE Networks 2026

3.2.11 Proxmox Configuration

1. Copy vNetC and SDLC vma files to host:

- Copy the the virtual machines (VMs) files to your Proxmox host. These are:
- `vzdump-qemu-vnc-x.x.x.x.vma.zst` : This is the backup for the **vNetC** VM.
- `vzdump-qemu-SDLC_x.x.x.x.vma.zst` : This is the backup for the **SDLC** VM.
- `vzdump-qemu-fabric-manager-monitoring.x.x.x.vma.zst` : This is the backup for the **Monitoring** VM.

2. List VMIDs in Use:

In Proxmox, each VM has a unique identifier called a VMID. To view the full list of VMIDs currently in use on your Proxmox server, run the following command:

`qm list`

3. Add the VMs and Assign Each an Unused VMID

After checking which VMIDs are already in use, select an *unused* VMID for each VM. Use the `qmrestore` command to restore the backups. For example:

- For vnetc, if VMID 110 is free, you restore it with: `qmrestore vzdump-qemu-vnc-6.3.0.3509.vma.zst 110`
- For SDLC, if VMID 111 is free, you restore it with: `qmrestore vzdump-qemu-SDLC_6.3.0.1.D20.vma.zst 111`
- For Monitoring, if VMID 112 is free, you restore it with: `qmrestore vzdump-qemu-fabric-manager-monitoring.1.2.18.vma.zst 112`

4. Edit the CPU, RAM, and Network for each VM:

Configure the resources for each VM via the Proxmox GUI (web interface).

- **CPU & RAM**: Set the correct number of CPUs and amount of RAM.
- **Network Configuration:**
- **vNetC:** For vNetC, configure the first network adapter (net0) to ensure it is connected to the correct bridge and has the correct VLAN Tag (if using VLAN trunking). Clear the MAC address field so Proxmox will generate a new one upon saving.
- **SDLC:** For SDLC, configure the second network adapter (net1) similarly, ensuring it connects to the correct bridge with the correct VLAN Tag. Clear the MAC address field so Proxmox will generate a new one upon saving.
- **Monitoring:** For Monitoring, configure the only network adapter (net0) to ensure it is connected to the correct bridge and has the correct VLAN Tag. Clear the MAC address field so Proxmox will generate a new one upon saving.

5. Start the VMs:

After configuring the VMs, start them from the Proxmox GUI: - Find the vnetc VM and click on "Start". - Find the SDLC VM and click on "Start". - Find the Monitoring VM nad click on "Start".

6. Continue Setup Process for vNetC, SDLC and Monitoring from Referenced Document(s):

After starting the VMs, follow the remaining instructions in the separate document(s) (linked below) to complete the setup for the vNetC, SDLC and Monitoring systems.

Key Steps Simplified:

1. Copy files to your Proxmox host.
2. List VMIDs to find available IDs.
3. Restore VMs using `qmrestore` with available VMIDs.
4. Edit CPU, RAM, and network settings for each VM.
5. Start the VMs.
6. Continue setup from the provided documentation.

Copyright BE Networks 2026

VNETC, SDLC AND MONITORING CONFIGURATION

- Configure the vNetC from the Console
- Configure the SDLC from the Console
- Configure Monitoring from the Console

Copyright BE Networks 2026

3.2.12 GNS3

We recommend running GNS3 on metal as opposed to a virtual environment. We have seen much better results.

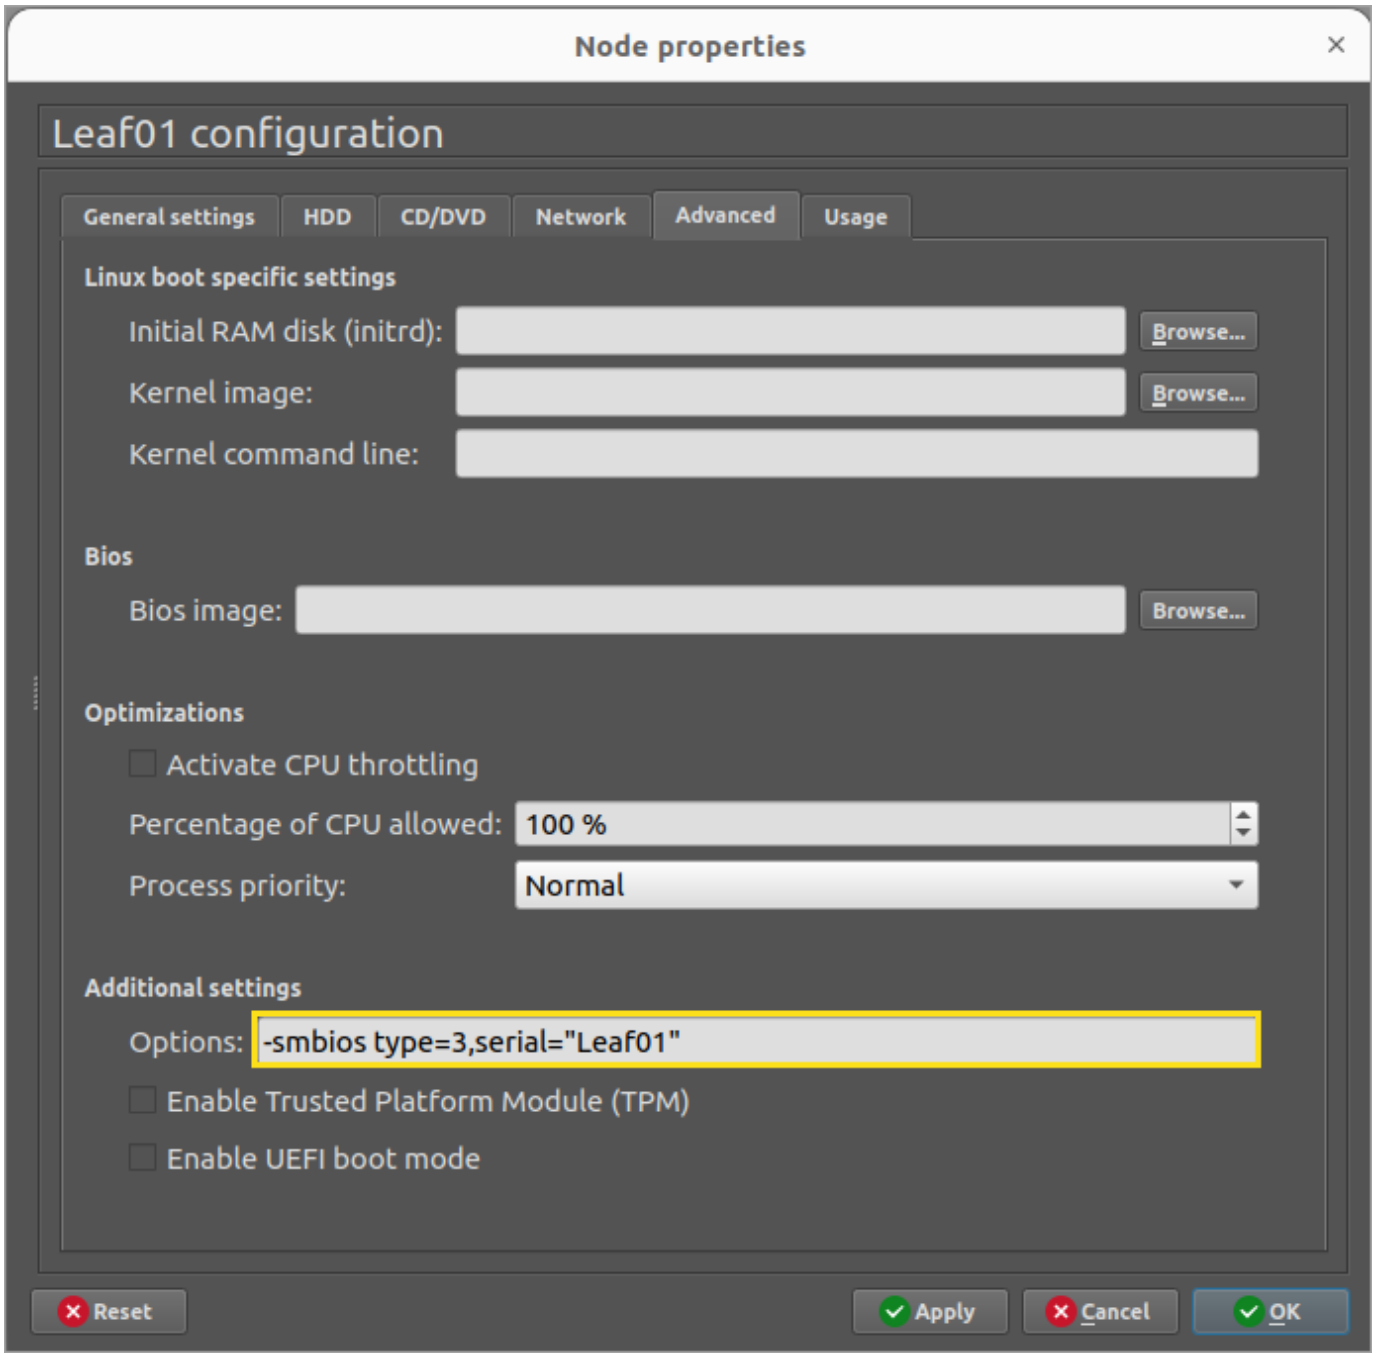
GNS3 Options to work with Verity

- Server Sizing to Support GNS3 with ~10 switches
- 88 cores, 200 GB Memory, 1 TB Disk

Each core

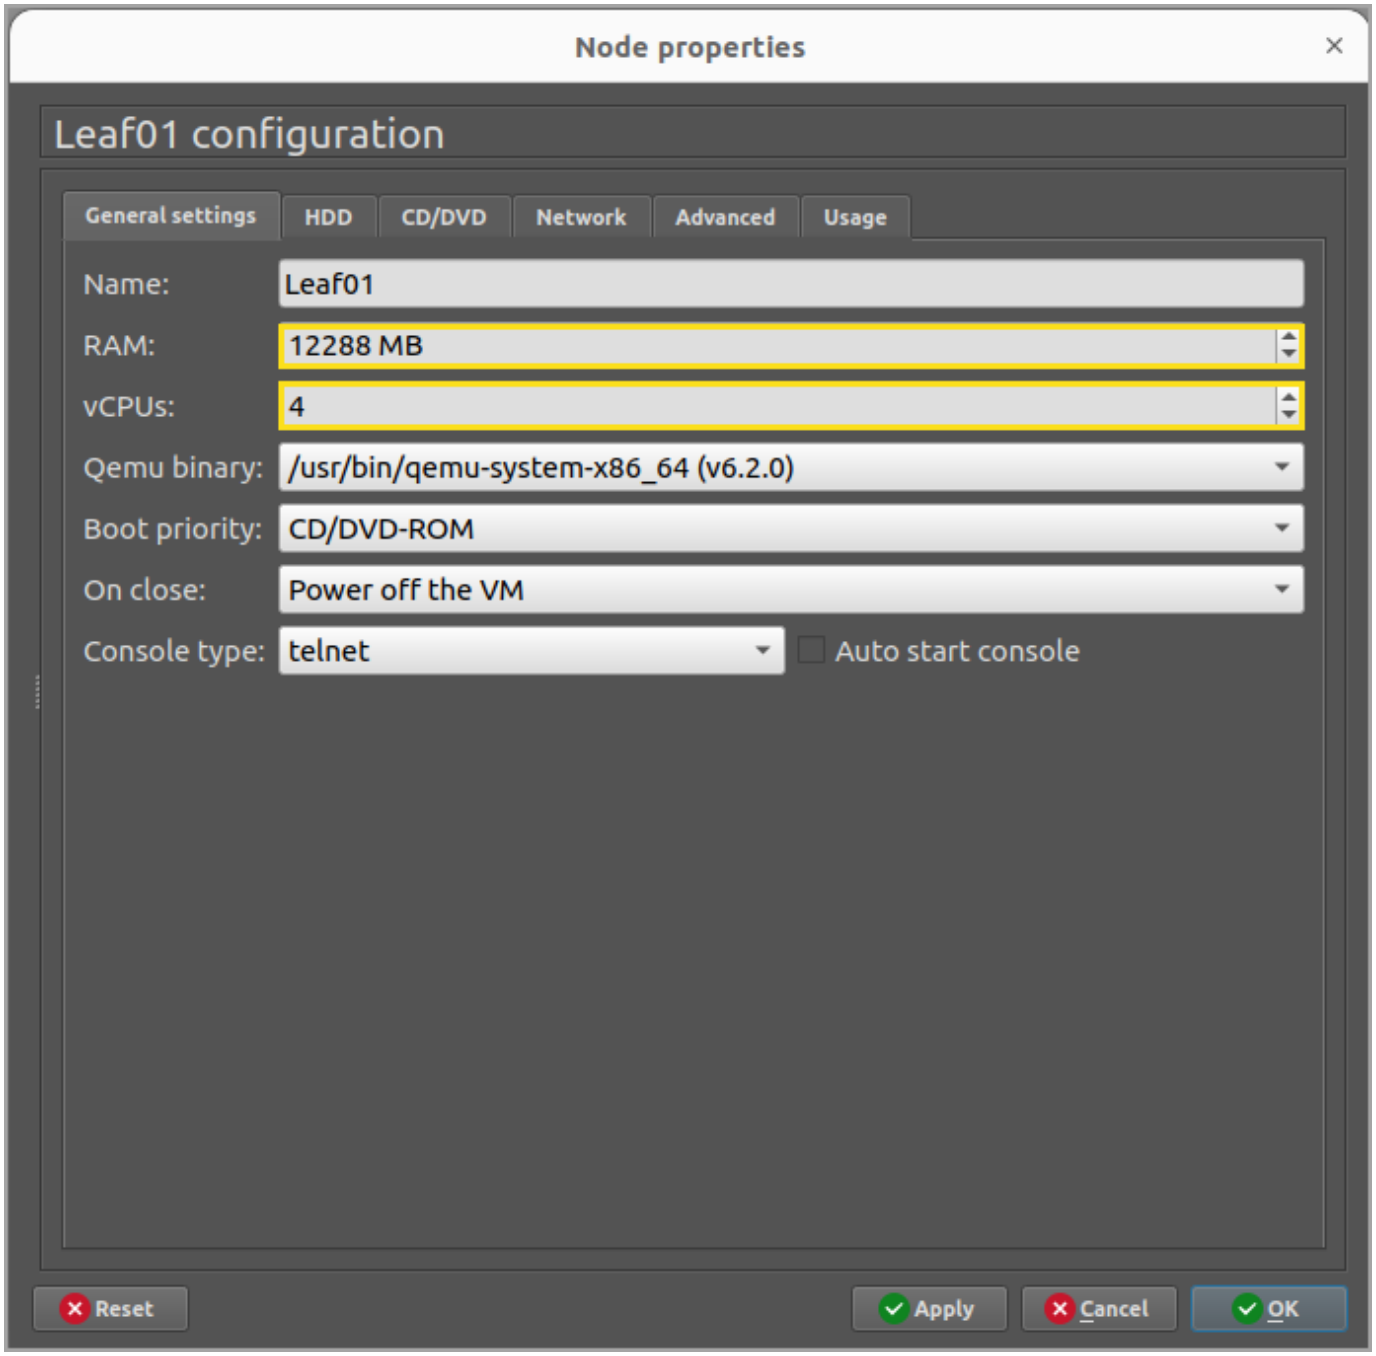
- Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz

In lieu of real hardware service tags/serial numbers, the string example below shows how we identify hardware in GNS3. This is critical for our device controllers and topology operations to work.



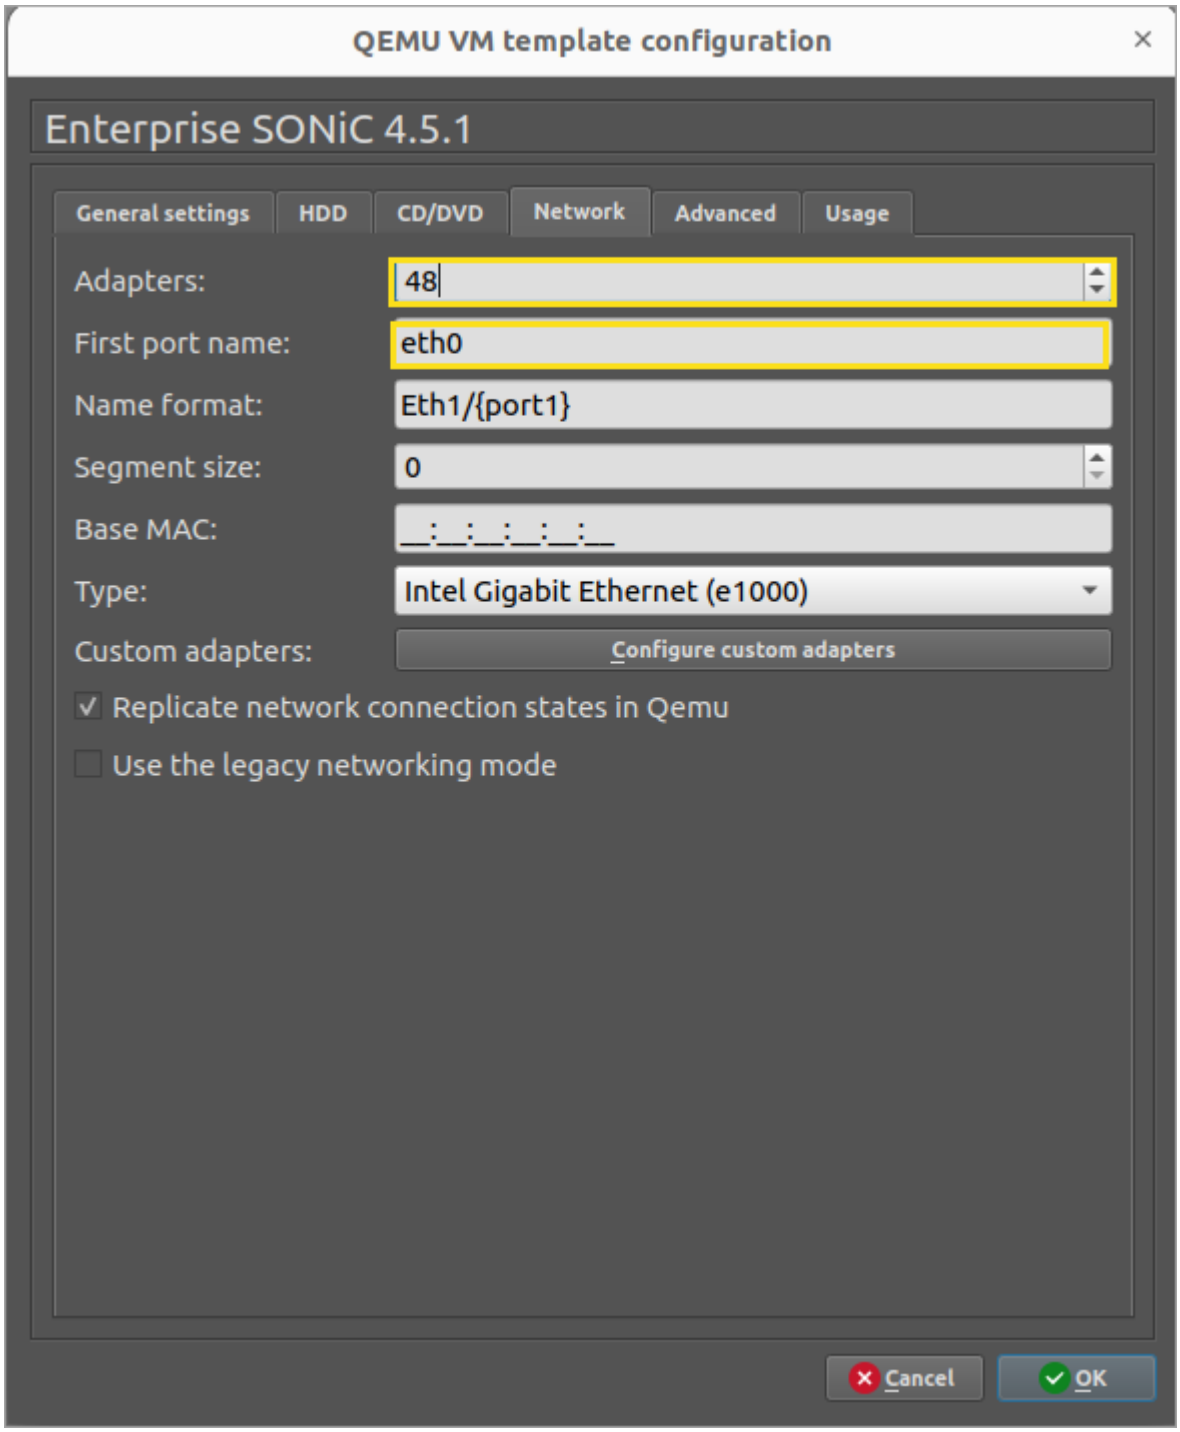
Copyright BE Networks 2026

Also, we have had the most reliable experience with the following RAM and CPU resources for Sonic simulated switches.



Set Default adapter count to 48

Copyright BE Networks 2026



Linux Network Recommendations

1. Set the MTU for all the interface being used to 9100. You can set that in **ubuntu netplan**. Run vi /etc/netplan/01-network-manager-all.yaml and add **mtu: 9100** to each interface. Below is an example of what it should look like:

```
#Let NetworkManager manage all devices on this system
network:
 version: 2
 renderer: networkd
 ethernets:
   eno1:
    dhcp4: yes
    mtu: 9100
 ethernets:
   eno2:
    dhcp4: no
    mtu: 9100
```


Copyright BE Networks 2026

```
vlans:
     eno2.70:
       id: 70
       link: eno2
       dhcp4: no
 ethernets:
 eno3:
   dhcp4: no
   mtu: 9100
ethernets:
  eno4:
    dhcp4: no
    mtu: 9100
```


Copyright BE Networks 2026

3.3 ONIE and ZTP

Requirements

The instructions in this document assume that you have installed the two Verity virtual machines required for use. It also assumes you have gone through the VM post installation process that includes the installation of license files and have created the management LAN subnet. Refer to the installation section for requisite procedure.

In this step, a DHCP server must be set up and the SONiC images as delivered by the vendor are installed into Verity.

The ZTP feature is a way to remotely set up switches without having to manually configure them on a hop-by-hop basis. The process consists of two steps: ONIE Boot process and ZTP Provisioning portion.

The ONIE process starts when a switch boots up from the factory without a SONiC image and it connects to a remote server to download the necessary software image. The actual ZTP process starts with a switch that has SONiC loaded and a default configuration. The Verity system creates a series of files that are used to prepare the SONiC switch with a minimal configuration to communicate with the Device controllers within Verity. DHCP options on the Verity SDLC server tells the switches the location of the remote server (vNetC) where the required files are stored. The system supports ZTP for both the management and spine/leaf switches.

3.3.1 Configuring DHCP

There are two options available for configuring DHCP:

- Use your own DHCP server - skip to the section titled **Using Your Own DHCP Server**.
- Install the internal **DHCP Server Application** by running the **Add a DHCP Server**.

Management Subnet

The address space for the management subnet is used for the vNetC (if in same subnet), SDLC and other internal system components used to manage the devices in the system as well as the devices themselves. This subnet, which was defined during the initial VM installation process, needs a section of addresses reserved for the static assignments of the system components as well as the management interfaces of the devices in the system. The dynamically allocated range is used to support the switches through the ZTP process as well as the Device Controllers that are created to communicate with the switches if they are set in DHCP mode.

If the managed devices are in a separate subnet, routed through external equipment (i.e. external to Verity), an external DHCP server is required in the subnet with the managed devices.

Option 1: Using the Internal DHCP Server

1. Open the Verity UI in your Chrome browser. Go to VNFs by double-clicking **Admin** icon and clicking the **VNFs** link.
2. Navigate to the section titled **System Applications** located within the VNFs section of the Verity dashboard and click the create button().
3. Click the **Add DHCP Server** button. 

DHCP Subnet Range

For the **DHCP Server**, setting the lease duration is required. The excluded address ranges are designated to reserve space for the static addresses assigned as described above. The addresses allocated for the Device Controllers and the switches during ZTP will come from within area between the excluded addresses. The subnet and this range should be sized accordingly based on how many devices are planned for the data center network. In this example, the first 100 addresses and the last 5 addresses of the management subnet have been excluded from the DHCP pool of addresses that can be allocated. 

Copyright BE Networks 2026

4. Check the box titled **Enable** and click the checkmark to save and start the internal DHCP service. 

Option 2: Using Your Own DHCP Server

1. It is not necessary to run the **Add a DHCP Server** installer as explained in the previous section, but please review the requirements described above for allocating IP address space. The following DHCP options are required which enables ONIE boot (Option 114) load and ZTP (Option 67).

Scope Setting	Value
Option 67	"http://**<vnetc-address or fqdn>**/download/ztp/file/ztp.json"
Option 114	"http://**<vnetc-address or fqdn>**/download/onie/file/onie-installer"

Note: If you are using secure (TLS) download option in the network, use "https" instead of "http"

3.3.2 Load SONiC Image

You are now going to load SONiC Image to VNetC. Before performing this action, you must obtain an image package from customer service containing the desired switch image files. You cannot load files directly supplied by 3rd parties in this step.

1. Go to the **Operations** page.
2. Click **NOS Images** in the **Version Control**.
3. Use the Browse Files/Drag-Drop interface to upload the pre-packaged OEM image file.
4. After upload is complete, click the **Deploy** (≫) button.

Congratulations

You have completed the Verity server setup and are now ready to begin importing network devices.

Copyright BE Networks 2026

3.4 Underlay Provisioning

Underlay Provisioning involved the following two steps:

1. Underlay Settings
2. Provision Reserved Ranges

3.4.1 Underlay Settings

Underlay Settings are available under the Fabrics navigation tab and are configured on a per-fabric basis. Underlay Settings are used to manage the time and conditions under which network devices detect, respond to, and stabilize after changes in network topology or connection states. In technical networking terms, these could be broadly described as "Network Convergence Timing Controls" - the settings that determine how fast and how precisely a network can detect, communicate, and stabilize after any type of network state change.

Underlay Configuration

Manages fundamental network infrastructure settings, including a shared Anycast MAC address for redundancy and load balancing. It controls how network devices handle address aging and connection restoration, ensuring stable and efficient baseline network communication.

Leaf Underlay Fabric Configuration

These settings control how leaf devices establish and maintain BGP connections, including timing for connection checks, routing updates, and connection establishment, ensuring stable and responsive network edge communications.

Spine Underlay Fabric Configuration

Defines how spine devices maintain and establish BGP (Border Gateway Protocol) connections. These settings manage network behaviors such as connection keepalive intervals, connection timeout periods, and routing information exchange.

For data center CLOS fabrics, Verity allows duplicate BGP AS numbers between spine endpoints in the same fabric. Superspine endpoints in the same fabric may also share the same BGP AS number.

Data Timer Configuration

Provides advanced timing controls that optimize connection stability and performance. These settings allow network administrators to fine-tune how devices handle interface state changes, link detection, and MAC address management. By adjusting Link State Timeout, Startup Delay, and MAC Holdtime, users can customize network resilience and reduce unnecessary connection disruptions.

Reserved Ranges

Read Reserved Ranges

Copyright BE Networks 2026

3.5 4. Onboarding Devices

3.5.1 Onboarding Devices

The most efficient method for onboarding devices is through bulk configuration using a system initialization file, referred to as an "FDC" file, as outlined in this document. Alternatively, devices can be individually onboarded through the Verity UI, though this method is best used when adding a few devices to a running system or replacing devices.

New Devices Starting Point

It is required that all SONiC devices being onboarded into Verity start in ONIE mode with no OS installed. This ensures that the device receives its load from the Partner Firmware Package installed on Verity and runs through the ZTP process.

Note

This section describes how to onboard devices in bulk using a system initialization file. This process creates Device Controller and Switchpoint objects and updates the Site object within Verity from the single import. To onboard individual devices see the section Device Management

System Initialization File

Device onboarding as well as the ZTP process, relies on an FDC file that defines the network devices and the management network, in a CSV file format. During installation, the user should prepare the file in Excel (or any CSV editor) and enter information into the designated fields. This is a list of the needed information to gather:,

- Service tag , serial number and chassis id of all switches (at least two of these values will be required, depending on the system's configuration and requirements).
- The hostname of the SDLC created during the VM installation process
- If using a Verity managed OOB management network, the designation of a management switch to function as your management Top of Rack (TOR) switch and the chosen uplink port you plan to use
- Leaf Pair Names if Multichassis LAG is being used
- Management network parametersbulk

Copyright BE Networks 2026

Example FDC files can be found in the Downloads section of the documentation.

1. Choose the CSV file built for the system. (Note: the file name must start with the string “FDC” and cannot have hyphen or slash characters in the name)

   Notes about the FDC File - The FDC file contains parameters used to build the data objects related to managed devices. Refer to the file to see examples of the entries related to the following column headers:

Column Name	Value
SWITCH NAME	This will become with switchpoint name in the user interface. Optionally is used to set the managed switch hostname.
ROLE	How the switch will be used.
FABRIC	Fabric.
SWITCH IP/MASK	Static IPv4 or IPv6 address installed in the switch by the ZTP process. For data center switchpoints, include a mask except for HGX switchpoints. Host-only masks are not allowed, such as IPv4 `/32` or IPv6 `/128`.
SWITCH GATEWAY	Static IPv4 or IPv6 gateway. This must use the same IP version as **SWITCH IP/MASK** and, when a mask is provided, must be in the switch IP subnet.
USERNAME	Switch username (should remain as "admin" for SONiC switches).
PASSWORD	Switch password.
POD	POD location of the switch.
RACK	Used for reporting and info, no action is taken on this parameter.
PAIRS	Name for each leaf pair paired when using MCLAG. Same name can only be used on two rows..
LLDP CHASSIS ID	Used as the basid for hardware detection using LLDP. Can be serial number, service tag or chassis ID. Using chassis ID is preferred, so that the "SWITCH NAME" will be programmed as the hostname in the device.
SDLC HOSTNAME	Controllers are containers inside the SDLC. A system may have multiple SDLCs.
IP SOURCE	optionally, device controllers can be set to static address. This entry specifies dhcp or static.
IP AND MASK	If using static addressing for controllers, this is the controller's address.
GATEWAY	If using static Device Controllers, this is the gateway for the controllers.
COM TYPE	Protocol or method a system uses to communicate with network devices.
ZTP IDENTIFICATION	The device identifier (serial number or service tag). Used by the system to create ZTP scripts for the individual device.
MGT TOR UPLINK PORT	Only needed to identify uplink port of the management switch at the top of the network.

Copyright BE Networks 2026

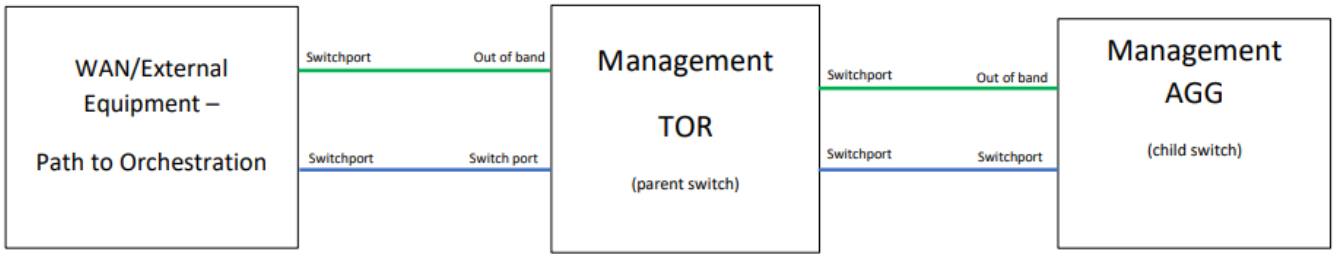
After the file is updated with the required information, import it by performing the following steps.

1. Click the **Operations** icon and then choose the **Import/Export** option.
2. Click **Import**. 
3. After importing the file, the system automatically creates the required provisioning objects for the system bring up.
4. Make sure that that the system read only mode is disabled (i.e., icon with white background vs. beige background). 
5. Click the **Operations** icon and then choose the **Import/Export** option.
6. Select **Import**. 
7. Choose the CSV file built for the system. (Note: the file name must start with the string “FDC" and cannot have hyphen or slash characters in the name.)
8. After importing the file, the system automatically creates the required provisioning objects for the system bringup.
9. Make sure that that the system read only mode is disabled (i.e., icon with white background vs. beige background). 

> **Device Management Protocol**
>
> The default protocol used for the device management in the automatically created Device Controller is "gnmi." Some switch models only support SNMP, CLI or other protocols. Please confirm with the hardware vendor's documentation to ensure that gnmi is supported for the models you are using.

Physical Connections

CONNECT MANAGEMENT SWITCH(ES)



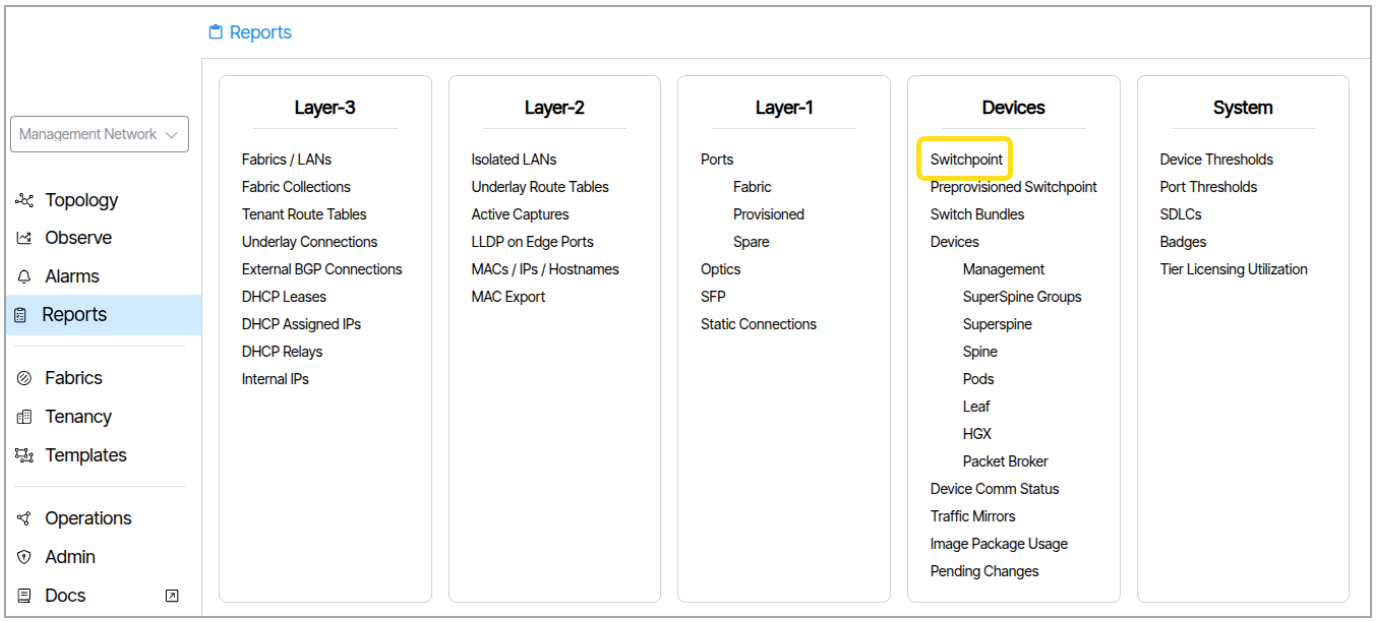
The ZTP process sets the new switch up based on the contents of the Device Controller object that was created by the FDC file import. As a default, the first 32 ports are configured to be possible uplinks.

ZTP requires two connections simultaneously from the switch. One of the first 32 switch ports is designated to be the uplink connection to the WAN/ Orchestration platform. The “out of band port” or “Management Port” of the switch is used to manage ONIE download and the ZTP process. Accommodations must be provided to allow the first switch (TOR) to go through the ZTP process. Once the TOR is up and running, connections for subsequent switches are made through the TOR. The management connections in the switching fabric are all untagged.

As shown in the diagram above, the management switches require two connections towards the installed applications to successfully complete the ZTP process. One connection is for the Management Switch “out of band management port” and the other connection is on a switched port. This switched port should be designated in the imported FDC CSV file only for the Management TOR switch. After connecting the device, the ZTP process will begin.

The next step is to “power up” the Device Controllers and set them to “Read write” mode. This is done through Reports. Open the Reports section of the main navigation UI and select the “Switchpoint” report.

Copyright BE Networks 2026



1. Select all switchpoints. 
2. Click the **Actions** button. 
3. Under **Switch** select **Read Only Mode Off**. 
4. Under **Switchpoint** select **Power Up** . 
5. Power up the physical switch.

   The ONIE boot load and ZTP process can take 10 to 20 minutes to complete. Upon successful completion and established communications with the controller, the management switch status transitions to green in the system UI. This means the device is now provisioned and the management switch is operational. To view the management switch go to **Topology/Management Network**

REMAINING MANAGEMENT SWITCHES

On the management TOR switch, physically remove the cable from the out of band management port. The same process with two cables is required for all management switches cascaded from the Management TOR

CONNECT SPINE AND LEAF SWITCHES

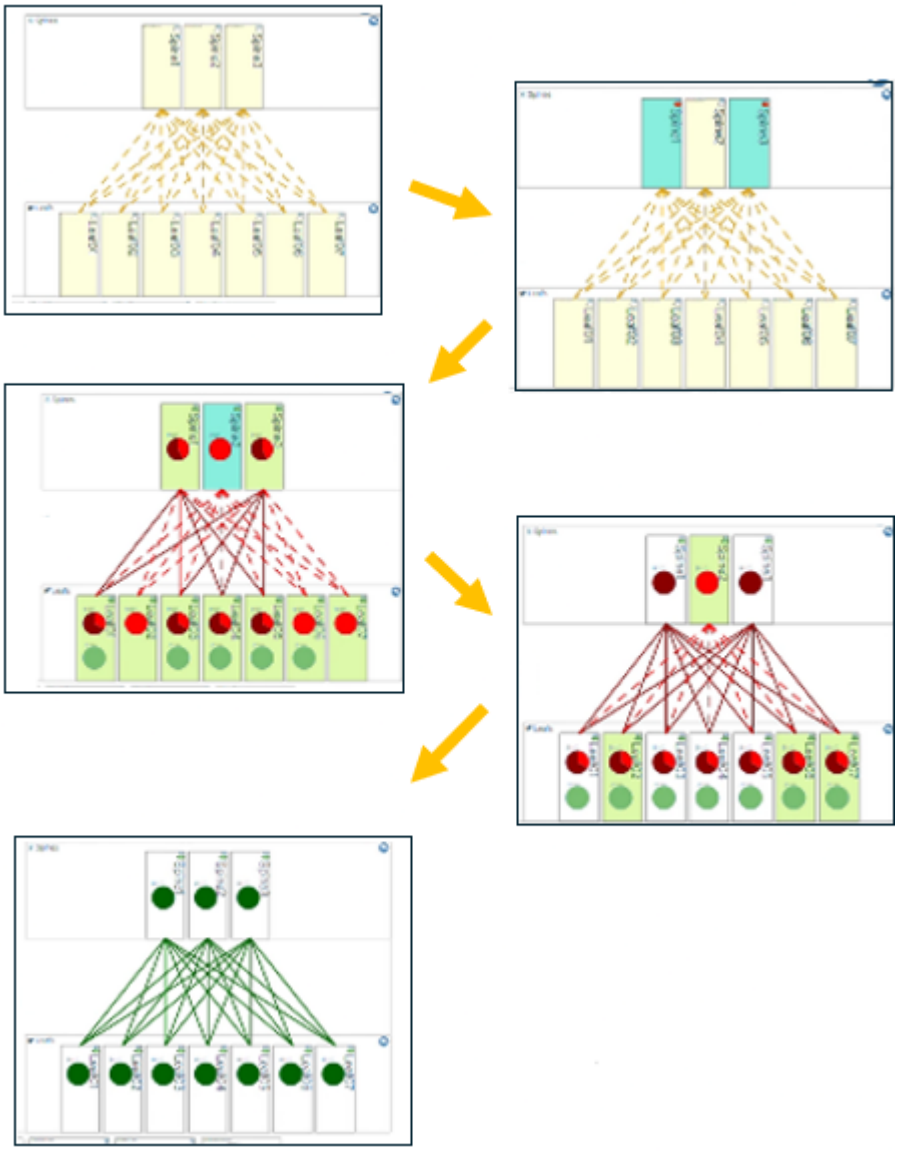
For spine and leaf switches, only the out of band management port is used.

1. Connect all other switches (super spine, spine, leaf) to their respective management switches. These switches remain connected to the management switch moving forward.
2. Connect the pairs of leaf switches together using any two switch ports you choose.
3. Connect the leaf switches to the spines and the spines to the super spines using any switchport.
4. Power up the switches.

Copyright BE Networks 2026

As you complete each step, the **Network View** updates in real time — the software discovers each switch and automatically configures the network underlay. As you perform the process the devices change color depending on their state.

Color	Meaning
Yellow	Switchpoint is Preprovisioned
Teal	Device is physically recognized, associated to the Switchpoint, and being registered
Green	Device is in provisioning process
White	Device provisioning is complete and is ready for use



Congratulations

You have now successfully installed Verity and activated the underlay network.

Copyright BE Networks 2026

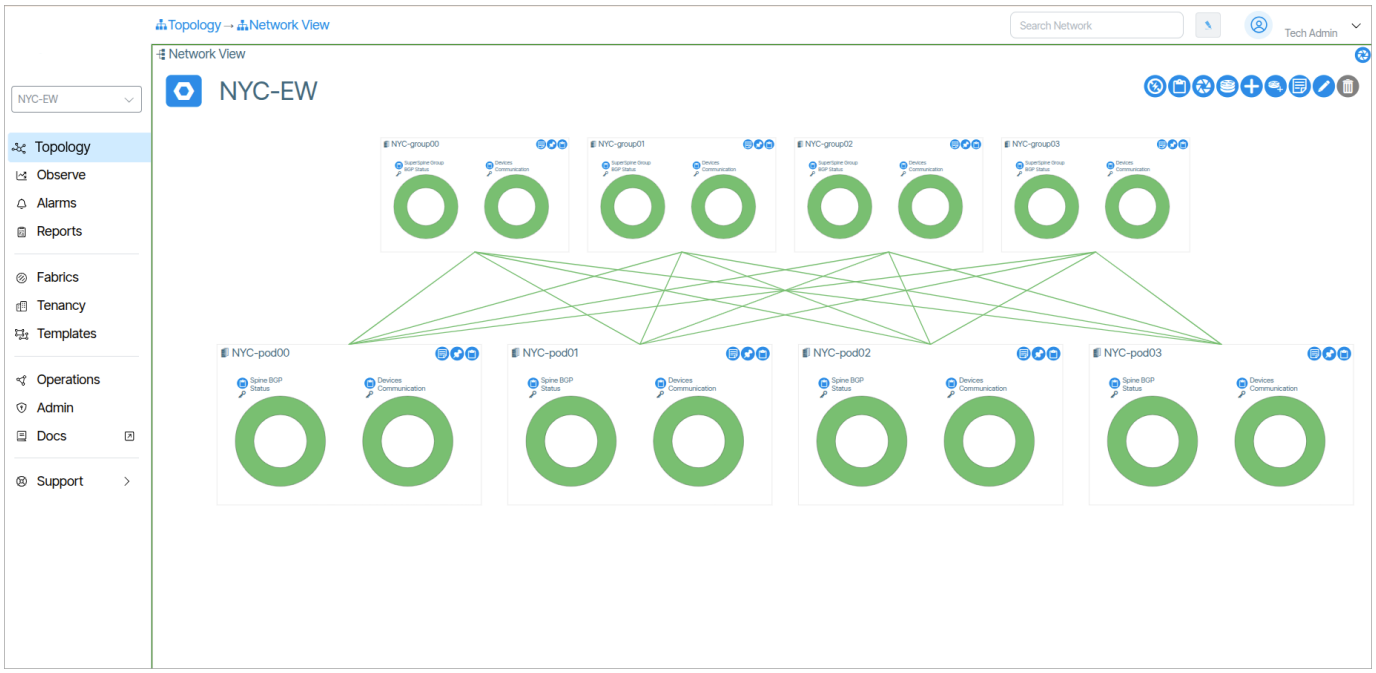
4. Using Verity

4.1 Navigation & UI

4.1.1 UI Navigation

Workspace

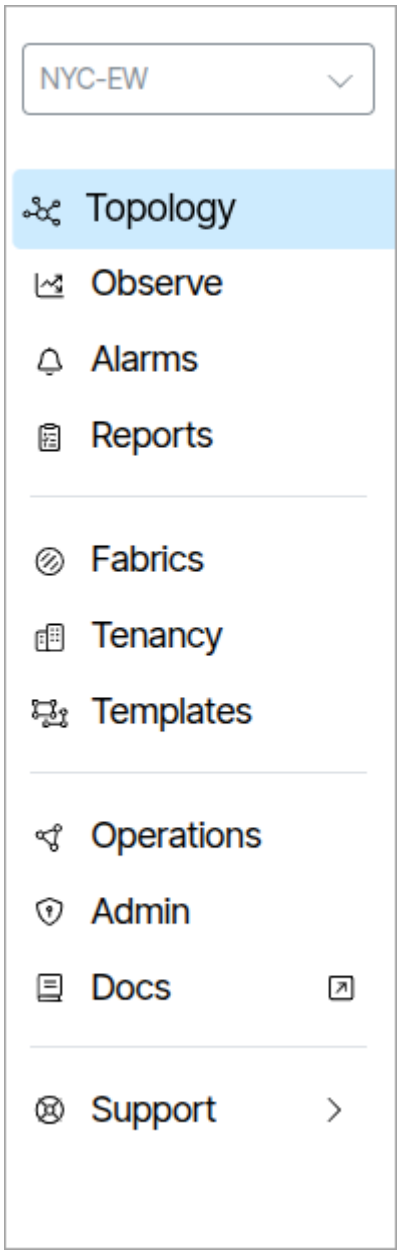
The **Workspace** is the *single pane of glass* through which you can configure the underlay and overlay networks, and apply these settings to network leaf ports. While the platform offers advanced capabilities, this portion of the documentation covers the essential navigation features to help you move efficiently through the user interface.



Main Navigation Bar

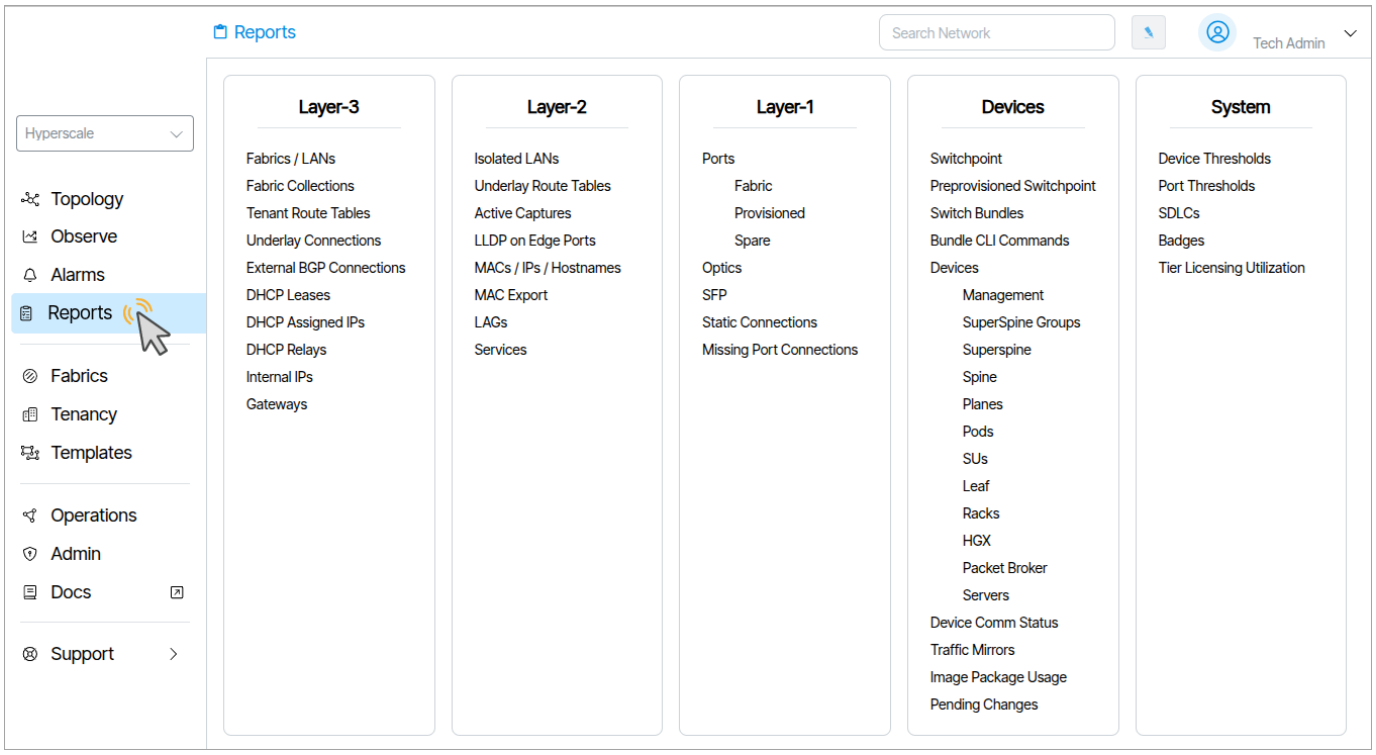
The **Main Navigation Bar** provides quick access to the different work areas.

Copyright BE Networks 2026



Double-clicking any item in the **Main Navigation Bar** reveals more options presented as categorized tiles. In the image below, the **Reports** navigation option is double-clicked, revealing a collection of new menu items.

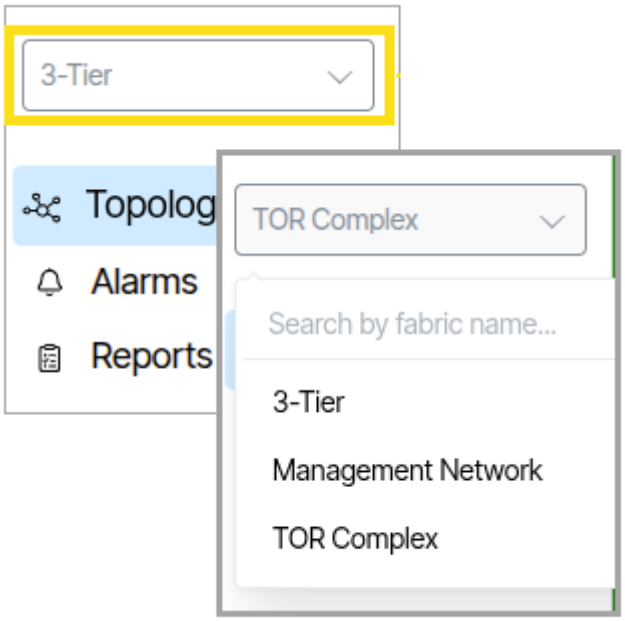
Copyright BE Networks 2026



Fabric Search

A **Fabric** is an isolated, self-contained arrangement of network interconnections that form a topology. This includes the switches, links, and logical structure. Multiple fabrics are administrable from the interface.

Fabric Search allows users to choose the Fabric to work with. When a Fabric is selected, the software interface updates to reflect the state and configuration of the Fabric throughout the entire application. For example, selecting a Fabric and immediately navigating to Network View displays the chosen Fabric's topology. Selecting Network View after choosing a different Fabric renders the topology specific to that selection.



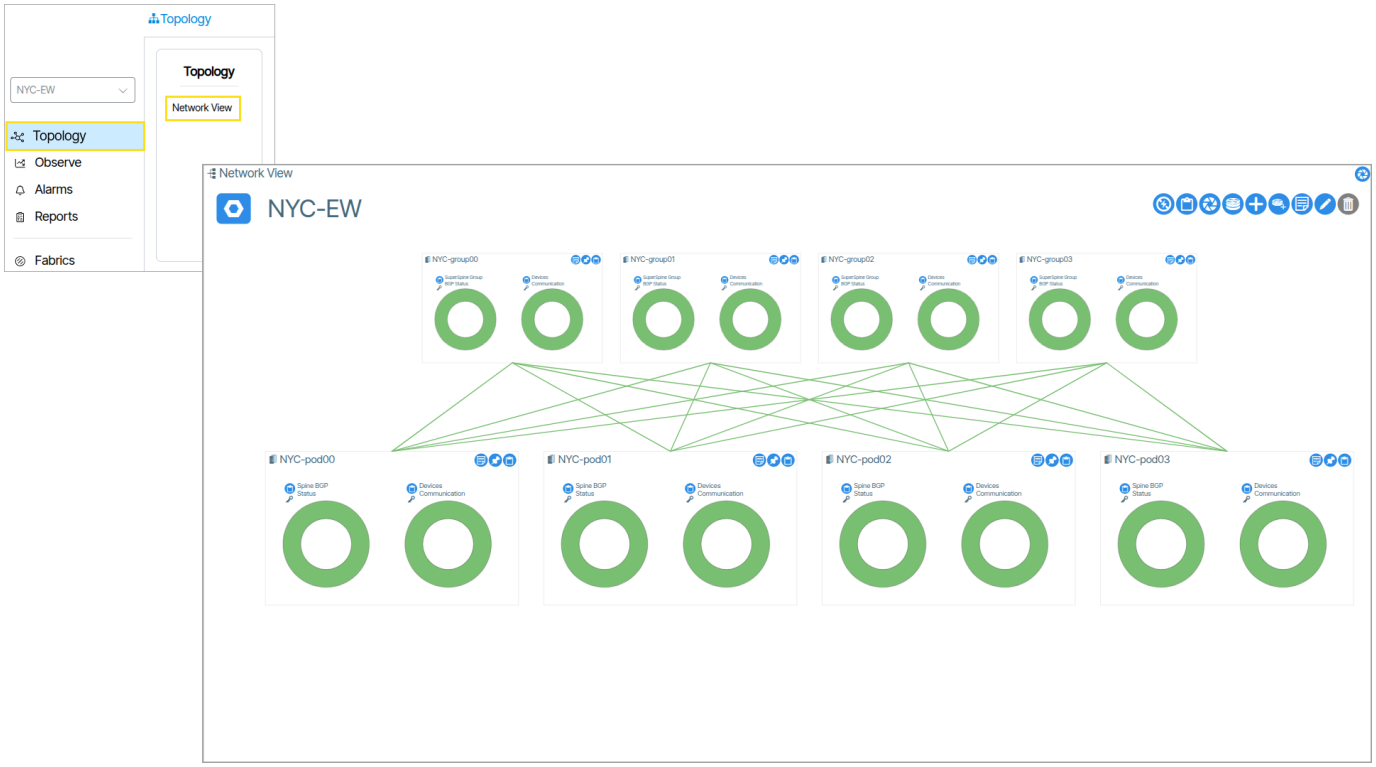
Copyright BE Networks 2026

Topology

The **Topology** menu provides visualizations of the network fabric and tools for managing physical connectivity. See Topology Overview for the full list of tools available.

NETWORK VIEW

Network View provides an interactive, zoomable visualization of your network topology. Use pan and zoom controls to navigate and explore devices and connections.

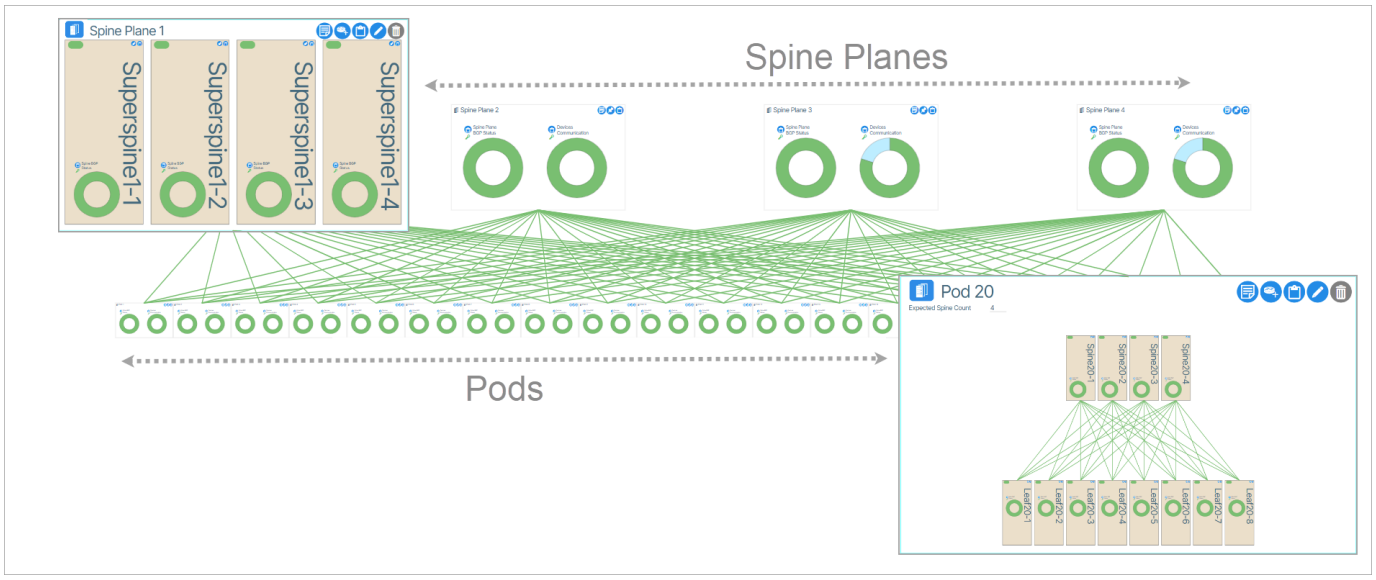


NETWORK REPRESENTATION

The devices are described in a Spine Plane Configuration

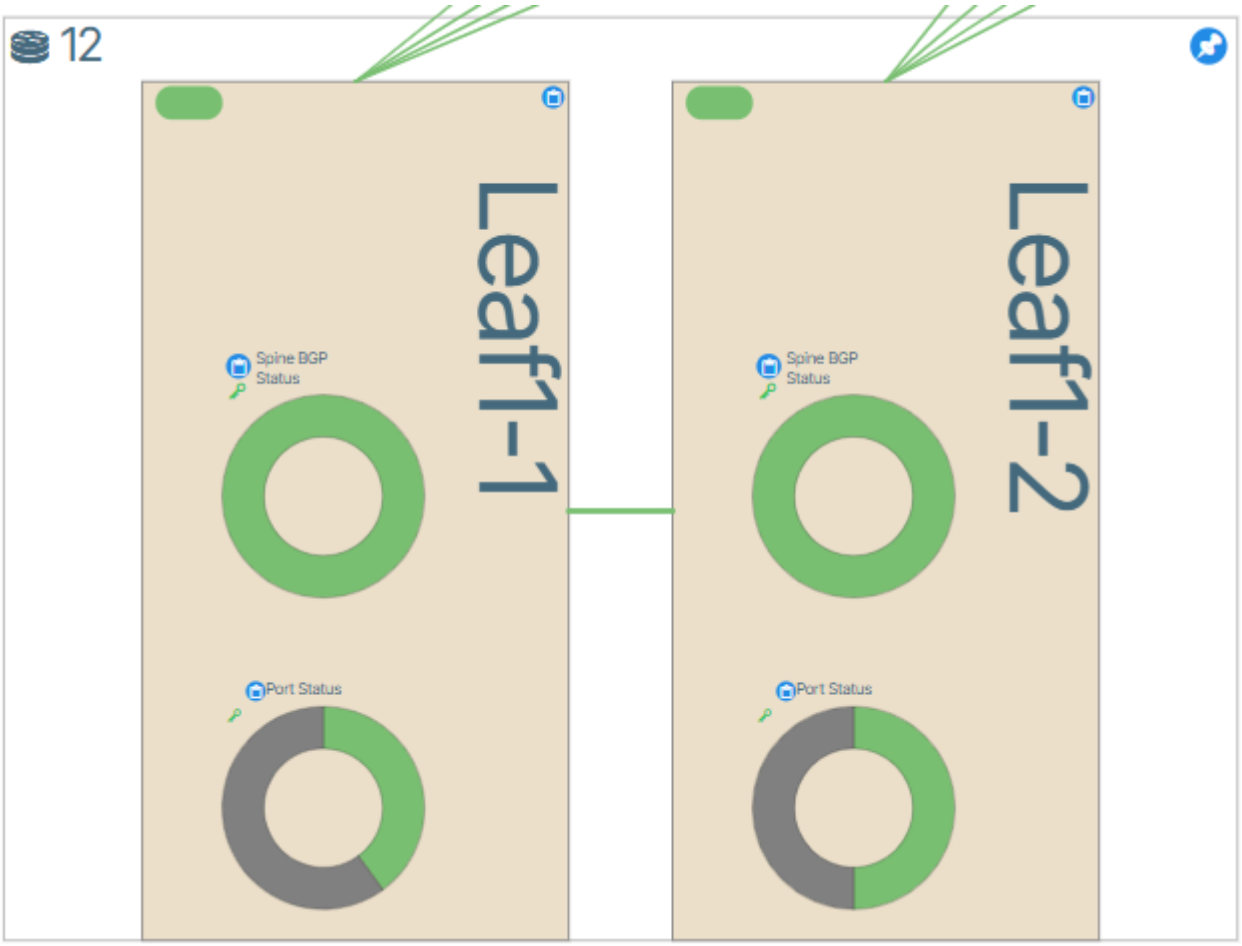
- PODs are composed of Spine and Leaf switches
- Spine Planes are composed of Superspines
- One Spine per POD connects upward to all Superspines in its assigned Spine Plane
- Multiple PODs can connect to the same Spine Plane, enabling inter-POD communication

Copyright BE Networks 2026



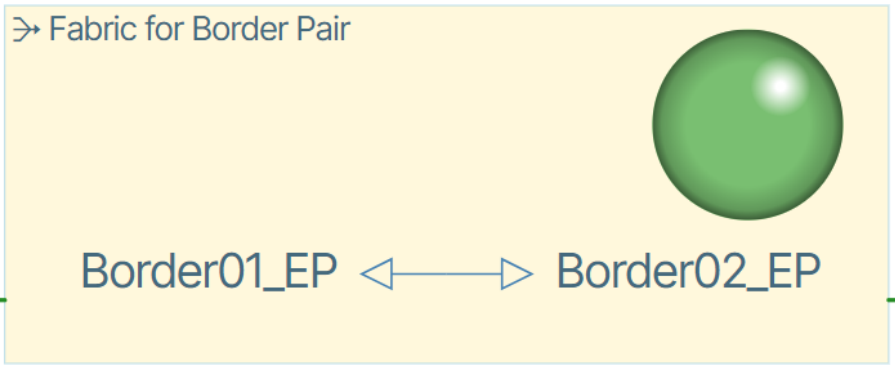
Leaf Pairs

Leaf Pairs are represented by a border encapsulating the paired switches.



When zoomed in, the fabric connection is represented by the following image.

Copyright BE Networks 2026



Color Connections

The state of devices is depicted by colored connection lines. The available state is depicted by a green color icon, while the down state is denoted by red, provisioning by gold and unstable by yellow.

Buttons

REGULAR BUTTONS

Regular buttons are scattered throughout the application and are used to perform different actions. If a button is not a Hex button it is a regular button. To perform an action with a regular button you left click it. There are many more regular buttons than those shown in the example below. A more detailed description of buttons is presented in the UI Icon Reference.



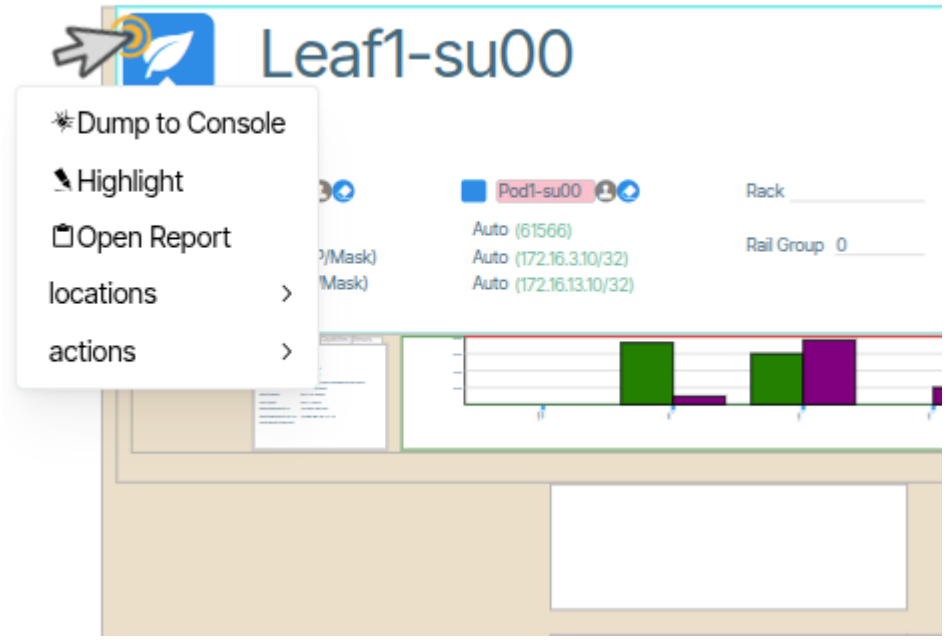
OBJECT ICON BUTTONS

Object Icon buttons represent tangible objects such as switches or other devices and respond to both right and left click mouse button actions.



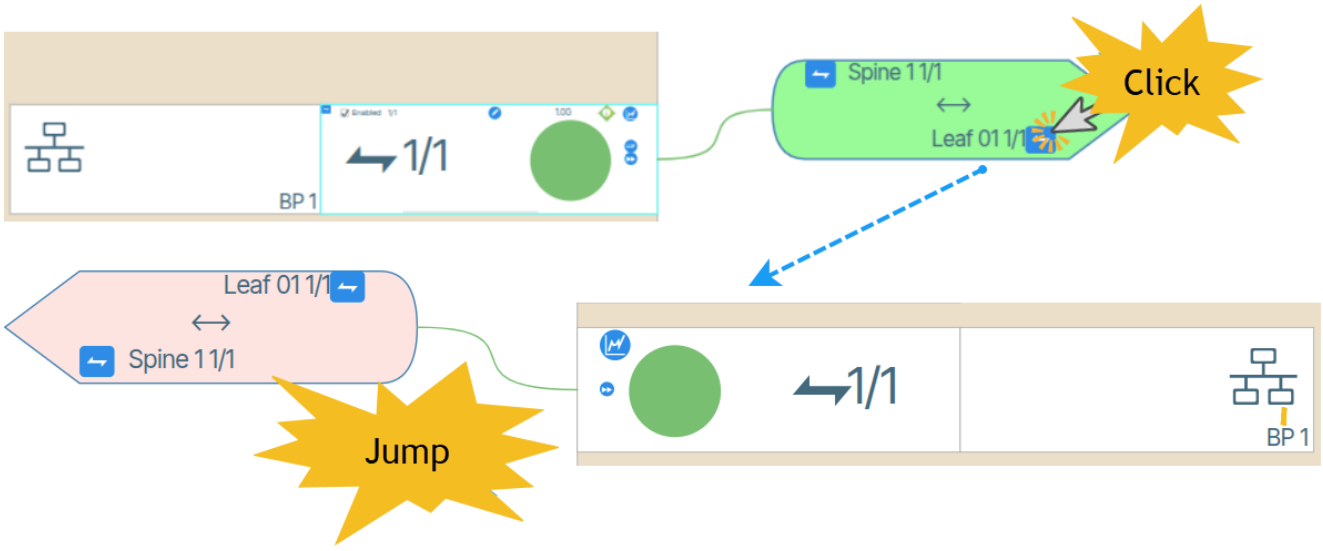
Right-clicking an Object icon button provides additional options that apply to the object.

Copyright BE Networks 2026



Depending on the virtual object, clicking a hex button forces one of two events to occur.

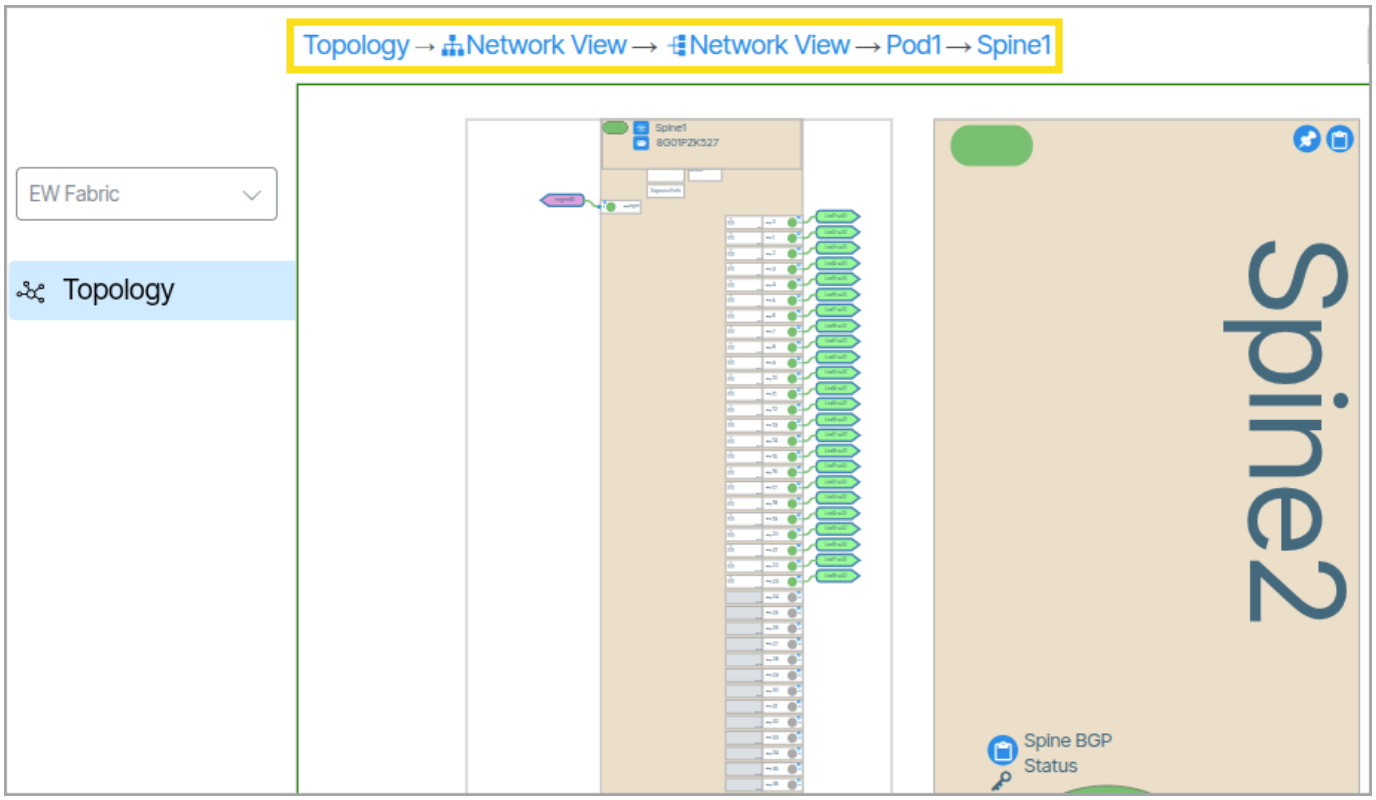
1. The click forces the window view to center on the Hex button container.
2. The click forces a window "jump" to the respective **Provisioning** section of the application. In the example below, clicking the object icon button named **Leaf 01 1/1** performs a jump to the corresponding device port in the configuration.



Navigation Depth Tracker

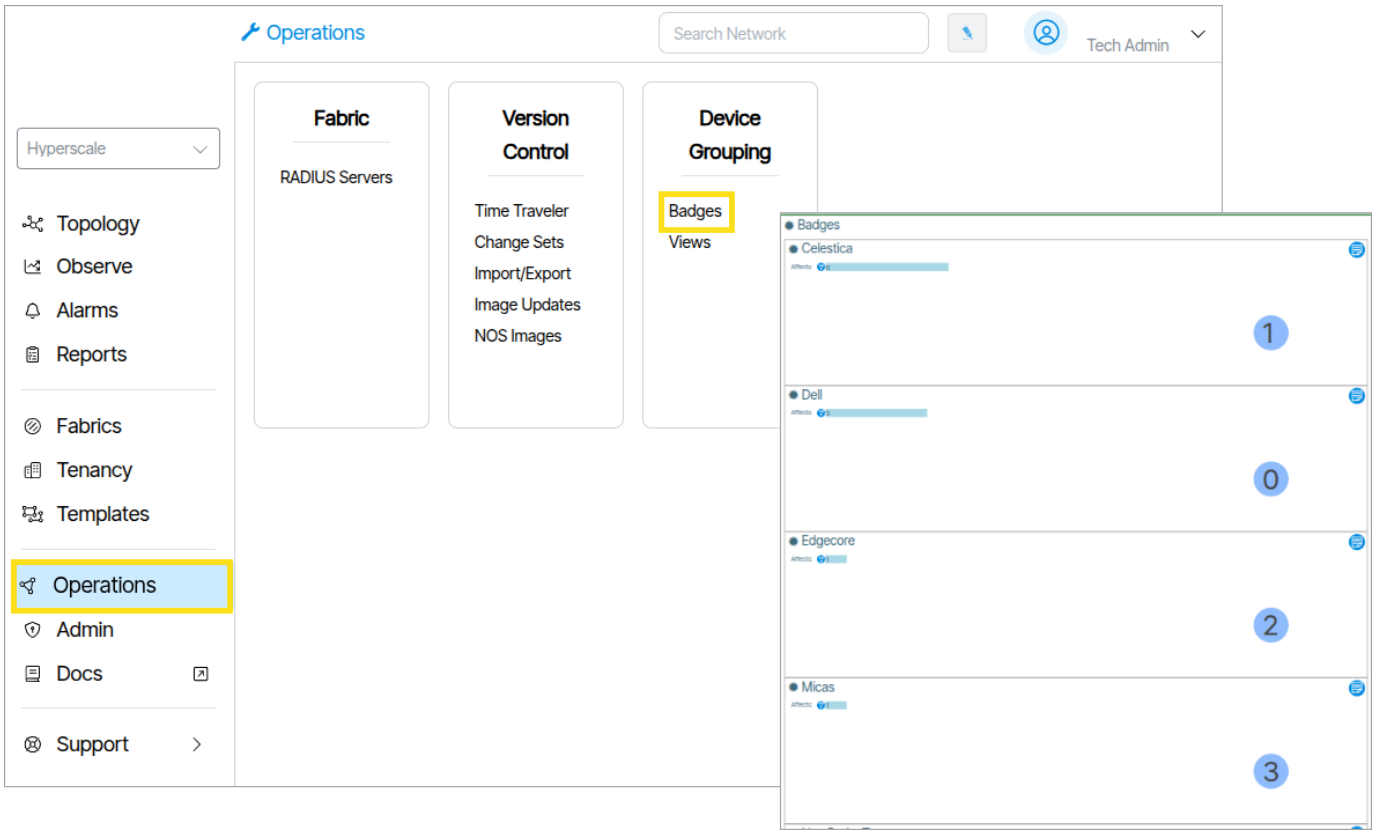
On the top left of the screen, there is a navigation depth tracker. This gives users a visual navigation path with clickable items used to return to earlier pages. An example is shown in the image below.

Copyright BE Networks 2026



BADGES

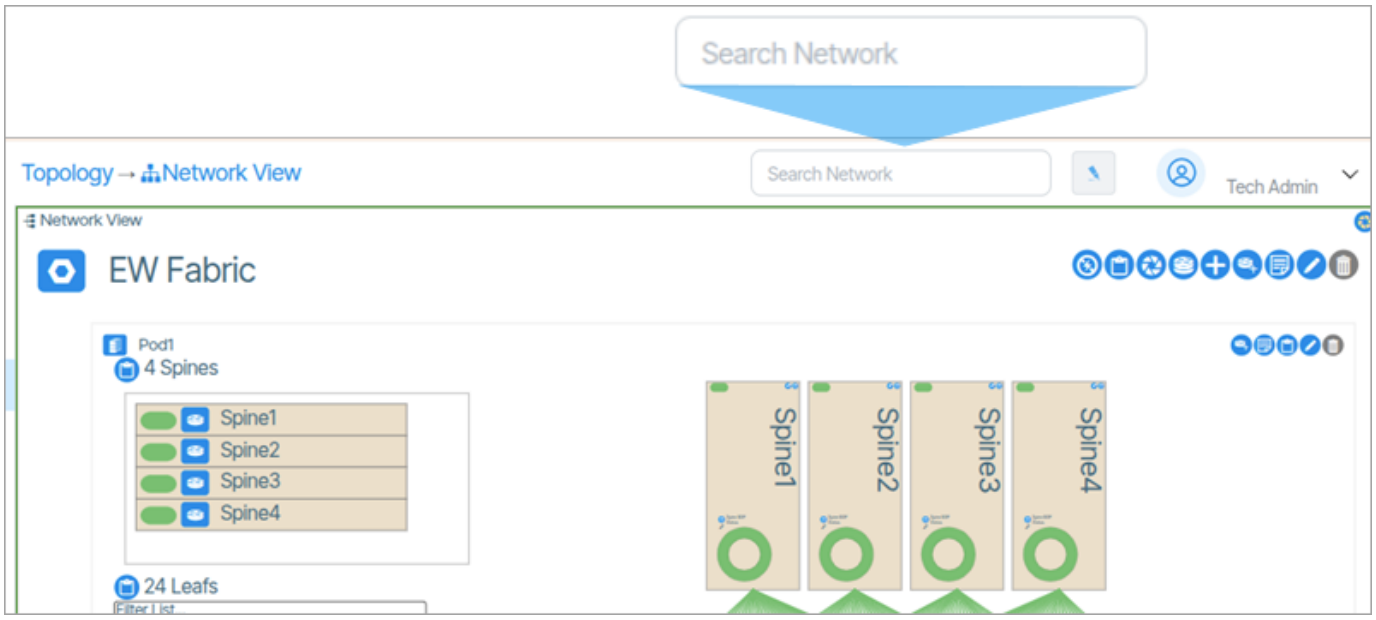
Badges are used to group devices visually.



SEARCH NETWORK

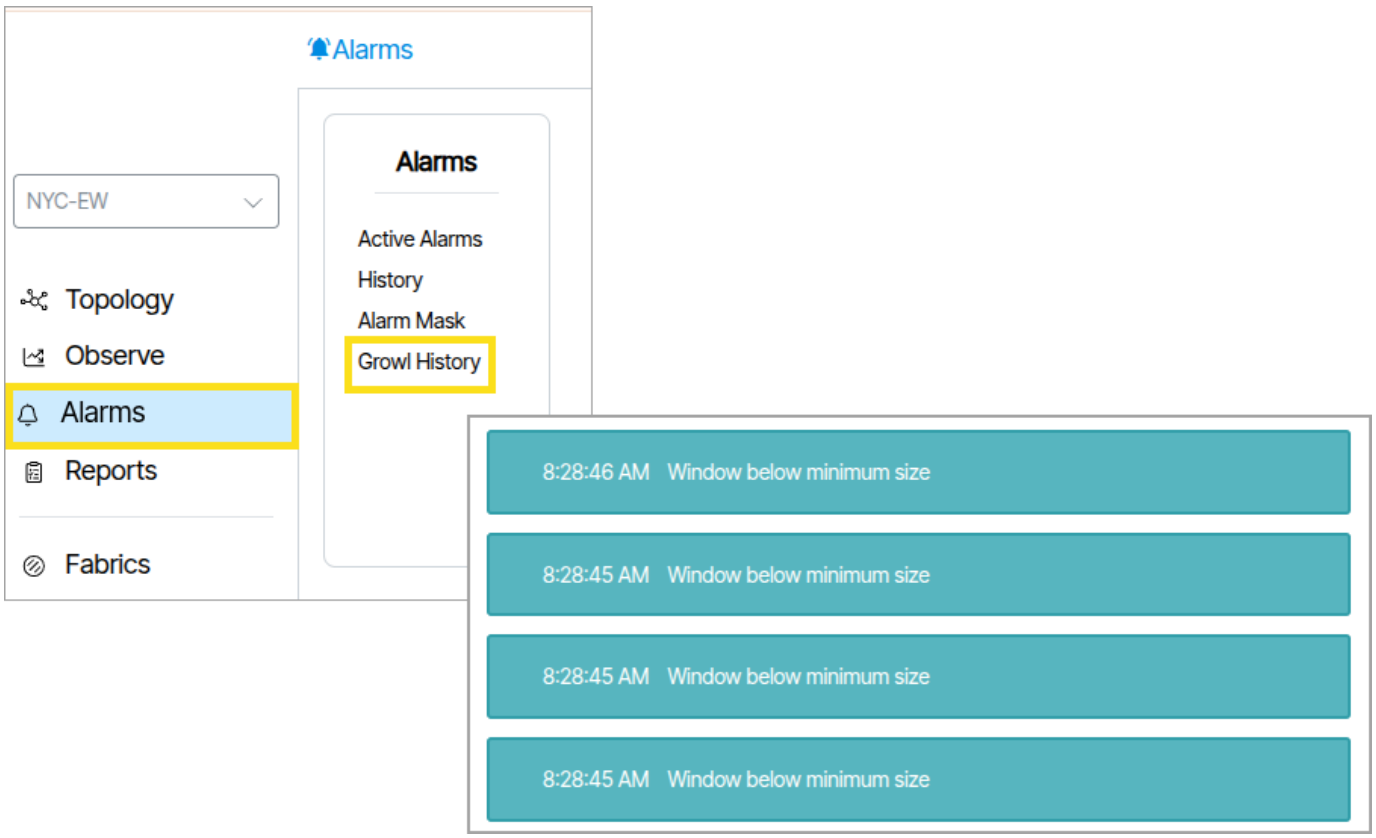
Search Network is used to search objects in the UI.

Copyright BE Networks 2026



GROWL HISTORY

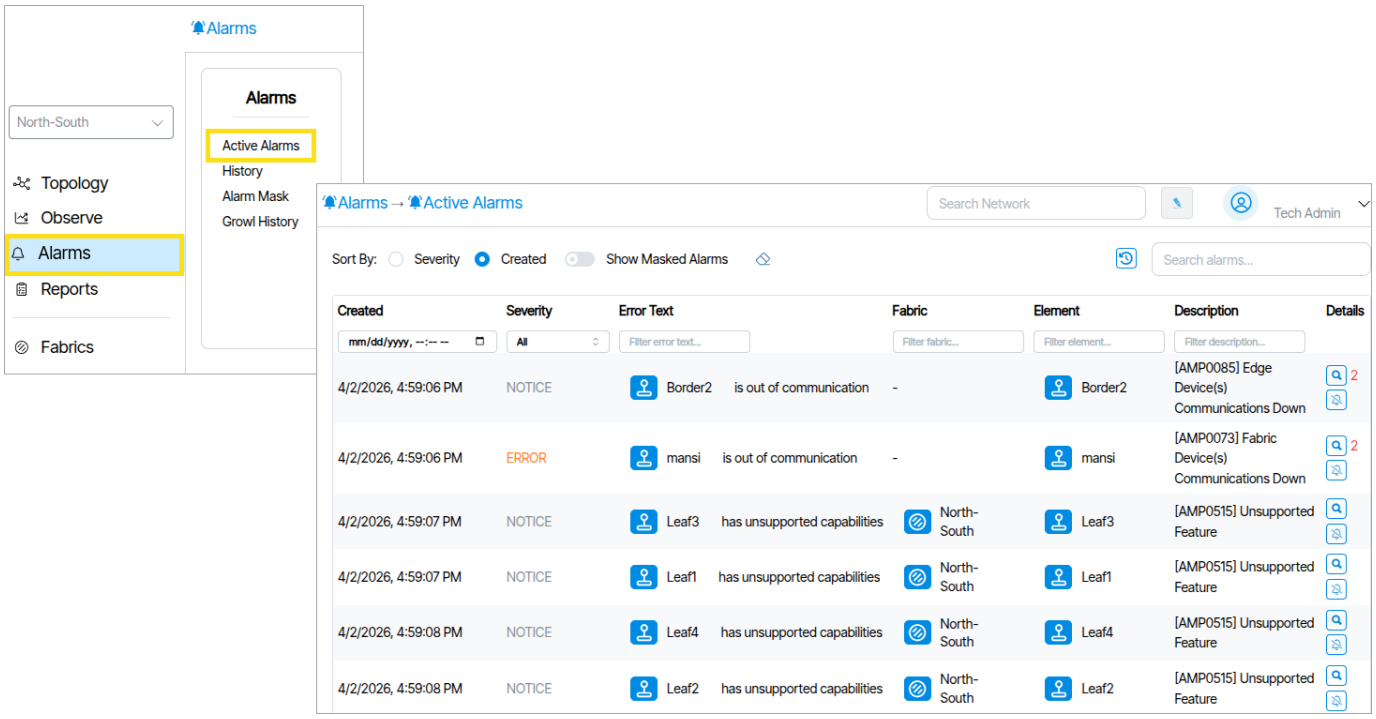
Growls are errors that occur between the client browser and the Verity orchestration platform. Holding the mouse over the growl gives the user more detail about the error. The Growl history is available at **Alarms/Growl History**.



ACTIVE ALARMS

Active Alarms are notifications that inform the administrator about changes to the network state. These notifications are displayed via the **Active Alarms** tab.

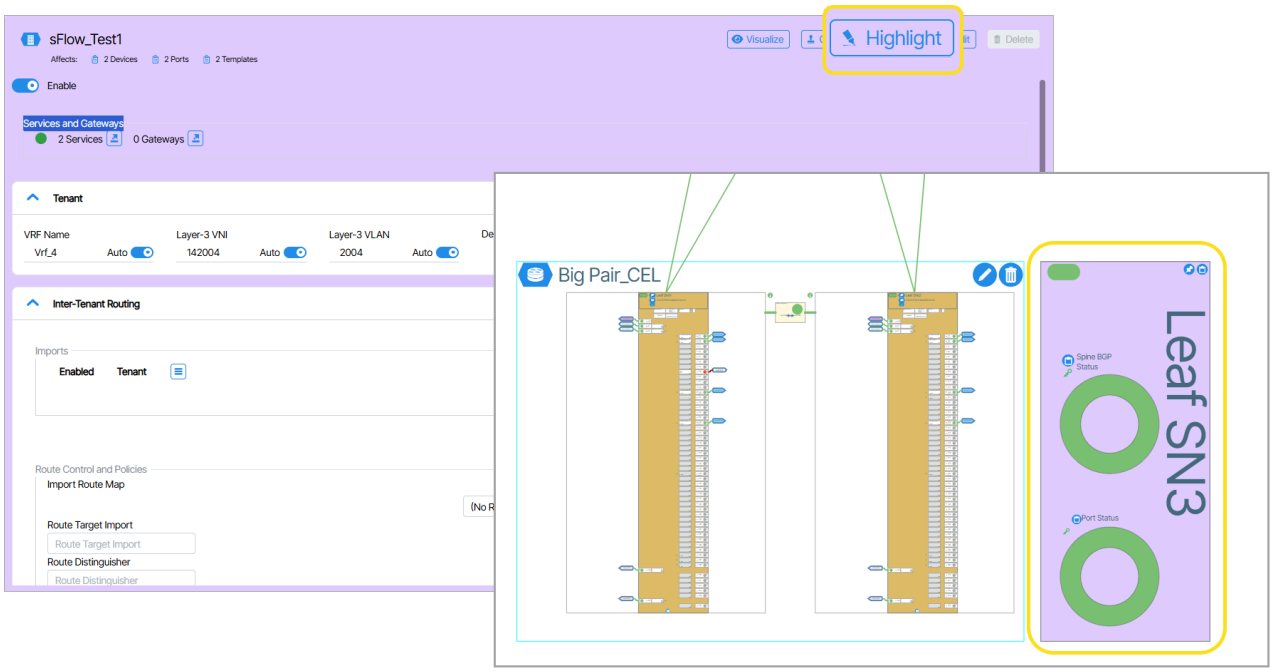
Copyright BE Networks 2026



HIGHLIGHTER



Highlighter is a tool used for tracing an object's data control flow across the UI. When you use the highlight tool, a purple hue is applied to the selected object and propagates to dependent objects in its control flow.

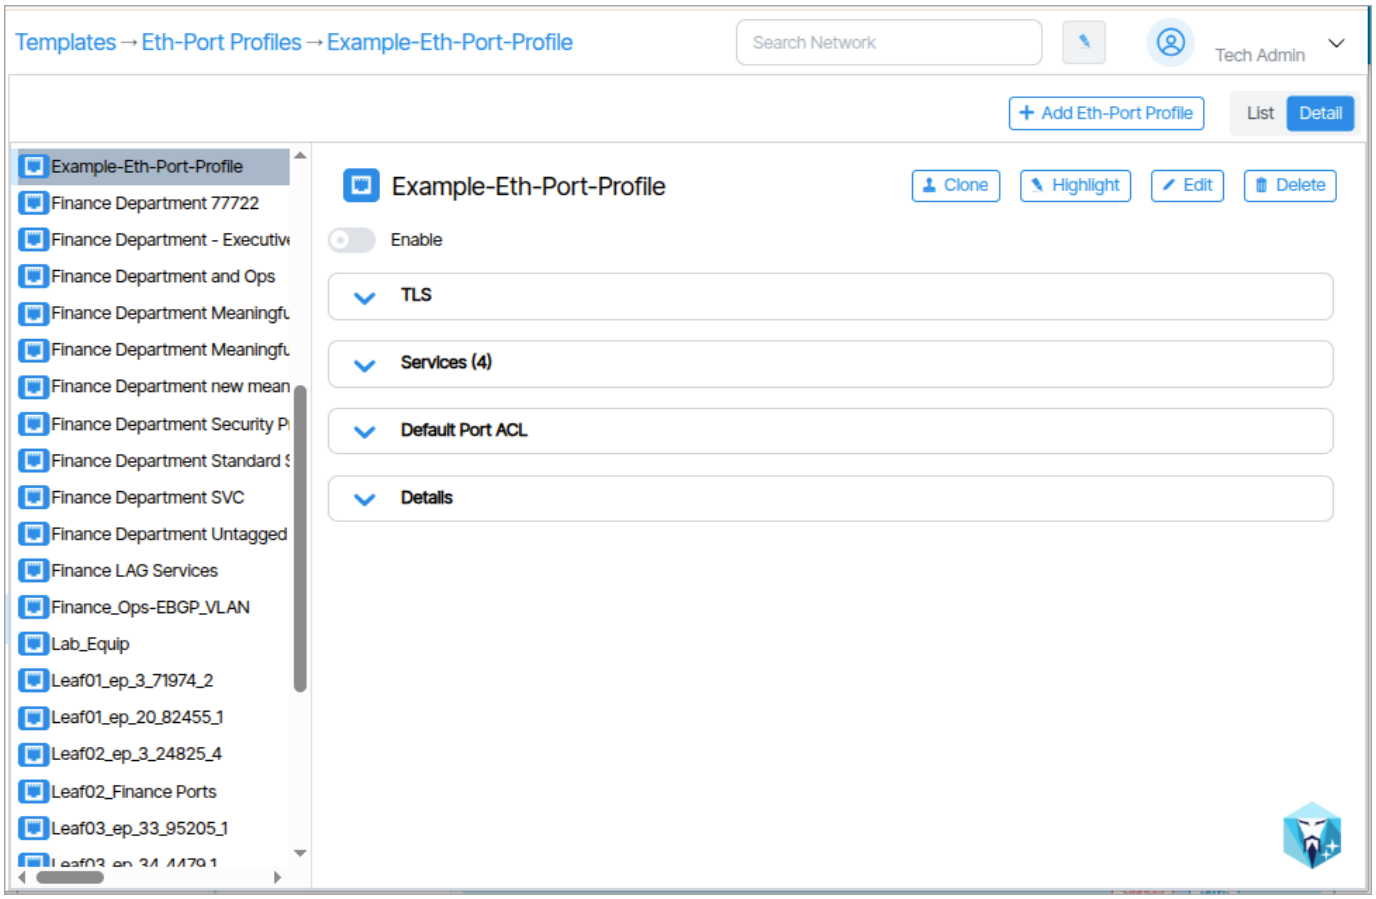


Copyright BE Networks 2026

4.1.2 Creating and Configuring Provisioning Objects

Provisioning objects are created using special configuration forms to define and apply settings. To create these objects, follow the steps below.

This is a general example that uses **Templates** and **Eth-Port-Profiles** for *demonstration*. Although the general creation and configuration process is the same, different provisioning objects will have their own unique parameters.



1. Click the link of the provisioning object you want to create. In this example (), **Eth-Port-Profile** is the provisioning object and is selected from **Templates**.
2. In the window that opens, click the **Add Eth-Port Profile** button. Depending on your selection, this button will contain text specific to the provisioning object such as **Add Service** or **Add Threshold** ().
3. In the field that appears, name the object and click the **Add** button ().
4. Fill in the parameters and settings of the object (). Fields are grouped together via an accordion interface. You can toggle the accordion open and closed by clicking its heading ().
5. The **Enable** switch is used to enable or disable the object ().

AFFECTS:

Affects: Indicates whether other provisioning objects in the system are linked to this object.

Copyright BE Networks 2026



ENABLE / DISABLE

Objects can be enabled or disabled by toggling the **Enable** button.

LIST / DETAIL

The window provides a visual choice between **Detail** and **List** view.

EDIT

To edit settings, click the **Edit** button.

DELETE

To delete the current object and all its settings, click the **Delete** button.

CLONE

To clone (copy) an object, click the **Clone** button. In the field that appears, type a new name for the clone and click **Create**. The name must be unique.

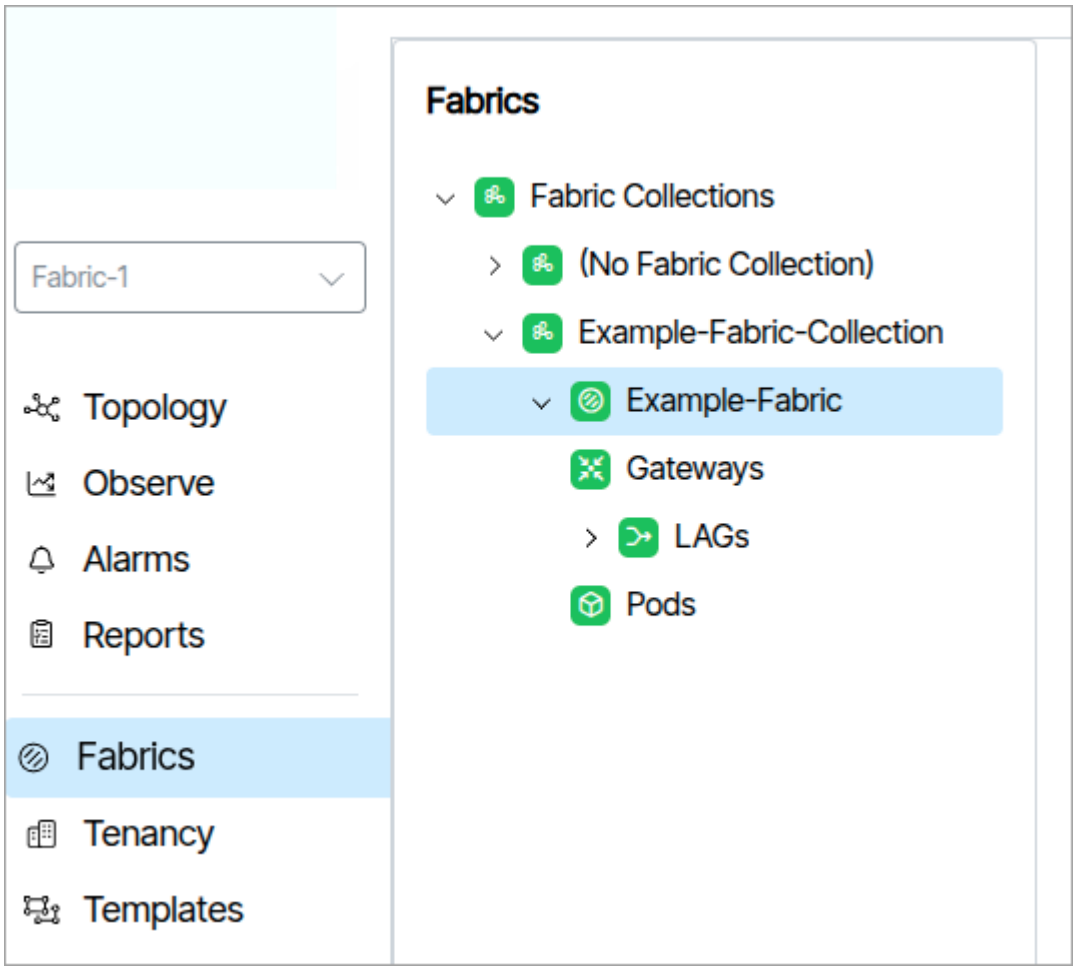
Copyright BE Networks 2026

4.2 Fabrics & Devices

4.2.1 Fabrics

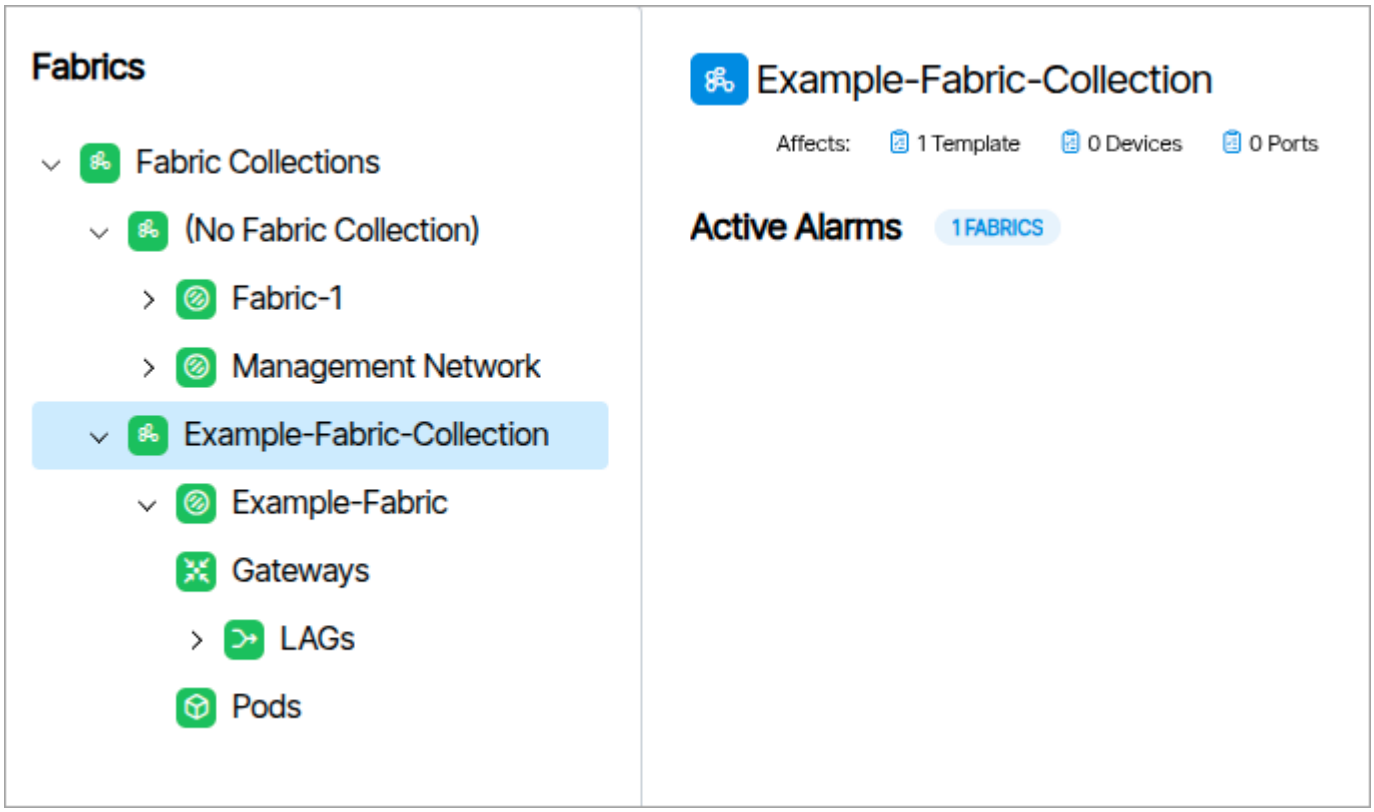
Fabrics are isolated, self-contained arrangements of network interconnections that form a topology, including the switches, links, and logical structure.

The **Fabrics** window gives administrators an overview of each Fabric by displaying the objects in a collapsible tree view.



Fabrics are organized into tree-node containers called **Fabric Collections**.

Copyright BE Networks 2026



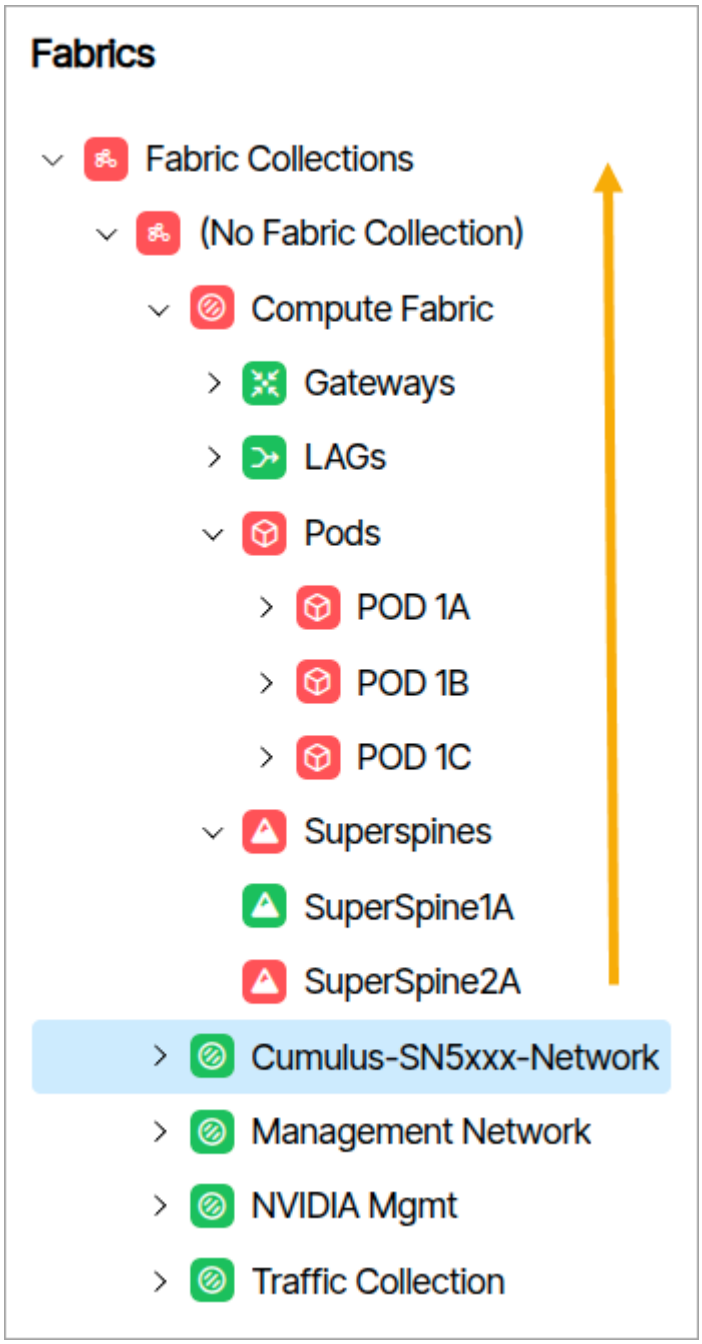
Colors

- The default color for a *disabled* fabric object is grey.
- The default color for an *enabled* working fabric object is green.

ALARM COLORS

When an object triggers an alarm, its icon changes to the alarm's color . The color change cascades up the hierarchy, always reflecting the most severe active alarm at each level. In the image below, the Fabric named Compute Fabric has child device(s) triggering a severe (red) alarm. This notification cascades through the entire hierarchy to the Fabric's parent.

Copyright BE Networks 2026



How to Create a Fabric Collection

1. Select **Fabric Collections**.
2. Click **Add Fabric Collection**.
3. In the form field that appears provide a name and click **Add Fabric Collection**.



Copyright BE Networks 2026

Creating Fabrics

1. Select the parent Fabric Collection.
2. Click the **Add Fabric** button. ()
3. A form appears requiring a name for the Fabric and a Fabric type. Fill in the fields and click **Add Fabric**.

The hierarchical object components display differently depending on the fabric type. The tree contains nodes that represent fabrics, categories for switch devices (LAGs, Pods, SuperSpines, Leafs, etc.), and the switch devices themselves.

Fabric Types

Management

A switch network providing connectivity to the out-of-band management ports of switches across the other fabrics.

- **Fabrics**
- **Management-Example**
- Switches

Data Center 3 Stage Clos

Spine-Leaf network configuration where traffic traverses three switch hops — leaf → spine → leaf — within a single POD.

- **Fabrics**
- **Data-Center-3-Example**
- Lags
- PODs
- Super Spines

Data Center 5 Stage Clos

Spine-Leaf network configuration traversing five hops — leaf → spine → superspine → spine → leaf — across multiple PODs. Superspines interconnect PODs for non-blocking any-to-any connectivity.

- **Fabrics**
- **Data-Center-3-Example**
- Lags
- PODs
- Superspines

Multiplane

- **Fabrics**
- **Multiplane-Example**
- Lags
- PODs
- Spine Planes

Copyright BE Networks 2026

Packet Broker

A Packet Broker is a network device placed between network traffic sources and monitoring tools (or other destinations).

- **Fabrics**
- **Packet-Broker-Example**
- Egress
- Ingress

SPECTRUM-X FABRIC TYPES

Editing Spectrum-X Capacity Settings

For Spectrum-X fabrics, the **Max SUs** and **SU Size** fields define the scale and capacity used for switch unit, leaf, spine, and HGX connection planning. These fields are editable only while the Fabric is empty.

After the Fabric has user-added content, Verity disables these fields to protect existing topology and connection calculations. Hover over the disabled field to view the reason, such as **Cannot change Max SUs when the Fabric has content** or **Cannot change SU size when the Fabric has content**.

To change these values, first empty the Fabric of user-added content. Auto-created empty planes and, for Spectrum-X 2 Tier fabrics, the single auto-created empty tier-2 POD do not prevent the Fabric from being treated as empty.

Spectrum-X 2 Tier and Spectrum-X 3 Tier fabrics can be deployed as single-plane or multi-plane systems. Each plane is an independent, complete network fabric. In multi-plane configurations, each SuperNIC's single high-speed port is split into multiple lower-speed links — one per plane — so the GPU sees full bandwidth while traffic is distributed across redundant, isolated fabrics for increased scale and resilience. In Topology planes are displayed as rectangular boxes  .

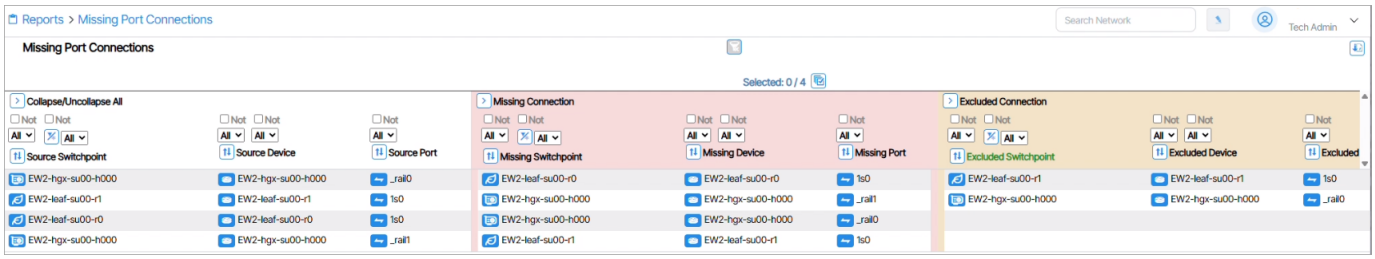
Spectrum X Cabling Validation

To verify that a Spectrum X system's physical topology is cabled correctly, complete the following steps:

- From **Reports**, select **Missing Port Connections**.

> **Note**
>
> If this **Missing Port Connections Report** is empty, it means the system did not find anything wrong with the cabling configuration.

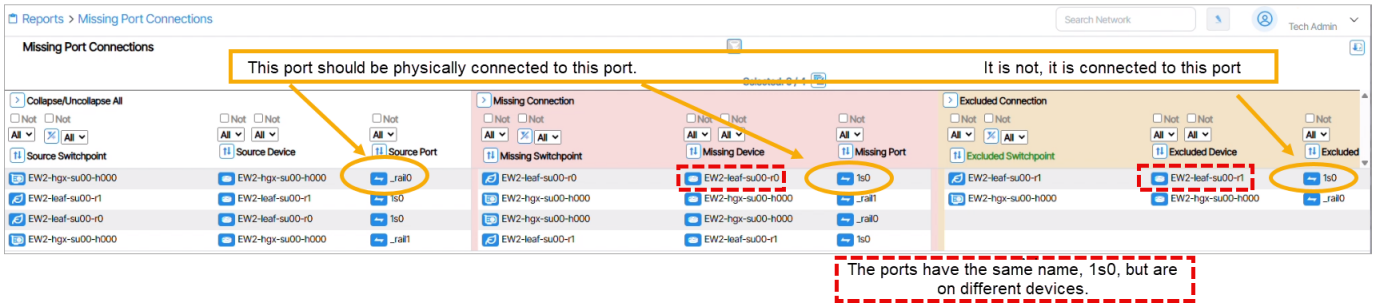
The information is presented as three columns. These are color coded as white, pink and peach. These columns have containing rows and each represents a physical connection point listed as a **Source Port**. The device and switch point associated with the **Source Port** are also listed as **Source Device** and **Source Switchpoint**. The first two columns (White and Peach) list pairs of devices that are disconnected from one another. A technician is expected to physically cable these devices according to the connection points presented. The third column (titled **Excluded Connections**) lists incorrect connections, meaning links that exist but are not part of the intended topology.A technician is expected to physically disconnect these connections.



Example

In the image below, port `_rail0` on device `EW2-hgx-su00-h000` should connect to port `1s0` on device `EW2-leaf-su00-r0`. That connection is missing. Instead, `_rail0` is incorrectly connected to port `1s0` on a different device, `EW2-leaf-su00-r1`.

Copyright BE Networks 2026



Spectrum-X 2 Tier

NVIDIA Spectrum-X Ethernet networking platform. Can be configured as a multi-plane system for increased redundancy and scale.

- **Fabrics**
- **Spectrum-X-2-Tier-Example**
- Spines
- SUs

Spectrum-X 3 Tier

NVIDIA Spectrum-X Ethernet networking platform. Can be configured as a multi-plane system for increased redundancy and scale.

- **Fabrics**
- **Data-Center-3-Example**
- PODs
- Superspine Groups

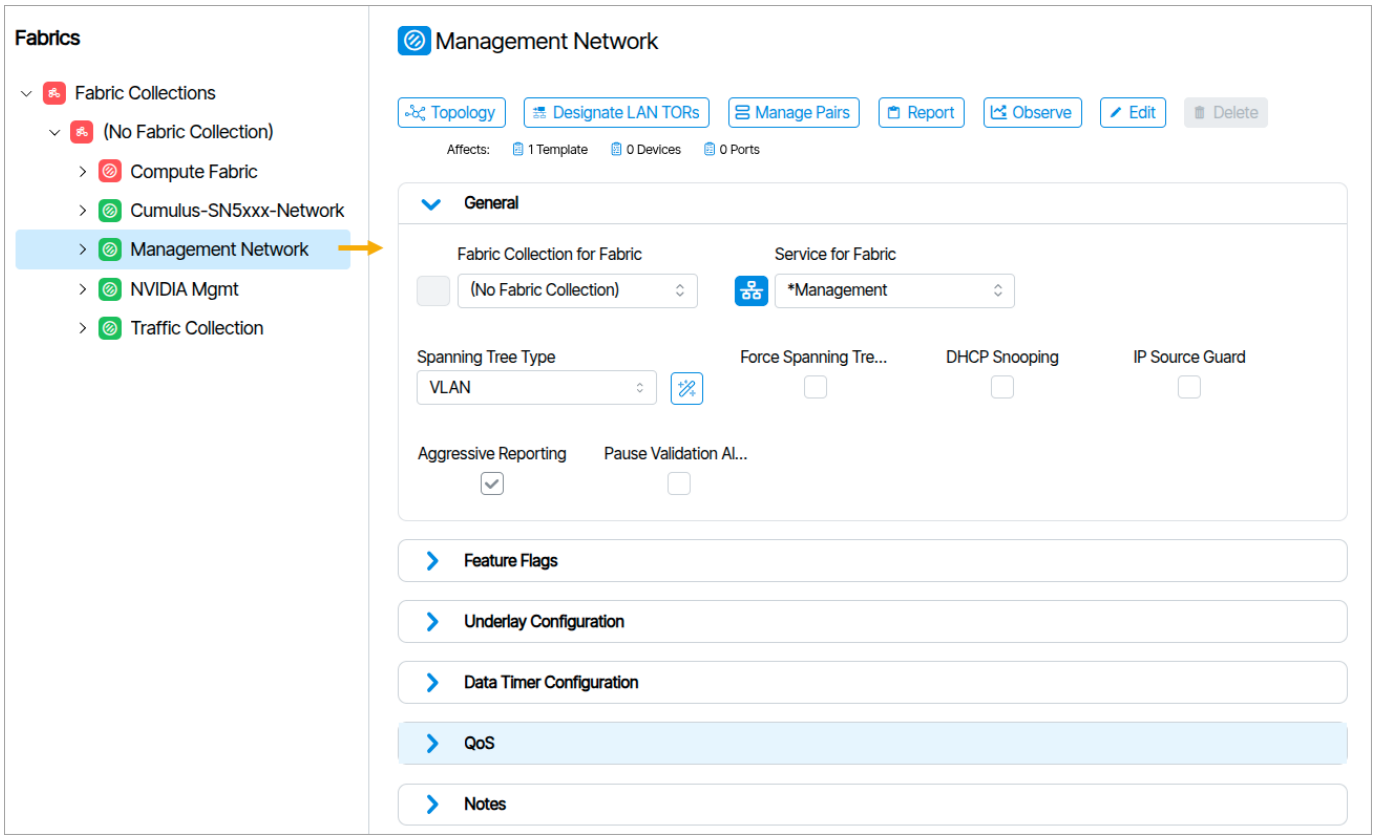
Inspector Window Object Details

The inspector window to the right reveals information and parameters of the selected tree object.

FABRIC DETAILS

Fabric objects have customizable fields that appear when you click them.

Copyright BE Networks 2026



Managed Pairs

The **Managed Pairs** button () opens a window that lets you manage and create Switch Pairs. Switch Pairs can be created as **Multi Chassis** redundancy pairs, or as non-redundant L3 leaf switch pairs when **Multi Chassis** is not selected.

Verity allows the pair to use leaf endpoints across PODs.

General

Fabric Type and other settings including DHCP Snooping.

Reserved Ranges

Provisioning reserved ranges specific to the fabric are listed here. Other Reserved Ranges are available in **Admin → System → Reserved Ranges**

Feature Flags

This section lists additional configuration options including settings to allow underlay connections between Pods, port admin and port status interval values.

Underlay Configuration

Manages fundamental network infrastructure settings, including a shared Anycast MAC address for redundancy and load balancing. It controls how network devices handle address aging and connection restoration, ensuring stable and efficient baseline network communication.

Leaf Underlay Fabric Configuration

These settings control how leaf devices establish and maintain BGP connections, including timing for connection checks, routing updates, and connection establishment, ensuring stable and responsive network edge communications.

Spine Underlay Fabric Configuration

For data center CLOS fabrics, spine switches in the same fabric may use the same BGP AS number. Verity allows duplicate BGP AS numbers between spine endpoints in the same CLOS fabric, including when the fabric-level Spine AS Number is not set. Superspine endpoints in the same fabric may also share the same BGP AS number.

Copyright BE Networks 2026

Defines how spine devices maintain and establish BGP (Border Gateway Protocol) connections. These settings manage network behaviors such as connection keepalive intervals, connection timeout periods, and routing information exchange.

Data Timer Configuration

Provides advanced timing controls that optimize connection stability and performance. These settings allow network administrators to fine-tune how devices handle interface state changes, link detection, and MAC address management. By adjusting Link State Timeout, Startup Delay, and MAC Holdtime, users can customize network resilience and reduce unnecessary connection disruptions.

Route Aggregation

Combines multiple specific IP routes into a single summarized route advertisement, so neighboring routers only need to know one prefix instead of many.

QoS

QoS is a set of techniques that prioritize certain types of network traffic over others when bandwidth is constrained or congestion occurs.

DSCP to TC Map

DSCP to TC Map is a translation table that converts quality-of-service markings between two different systems. This feature maps incoming DSCP priority markings (0-63) to internal Traffic Class queues (0-7) to ensure high-priority traffic like voice and video receives proper handling.

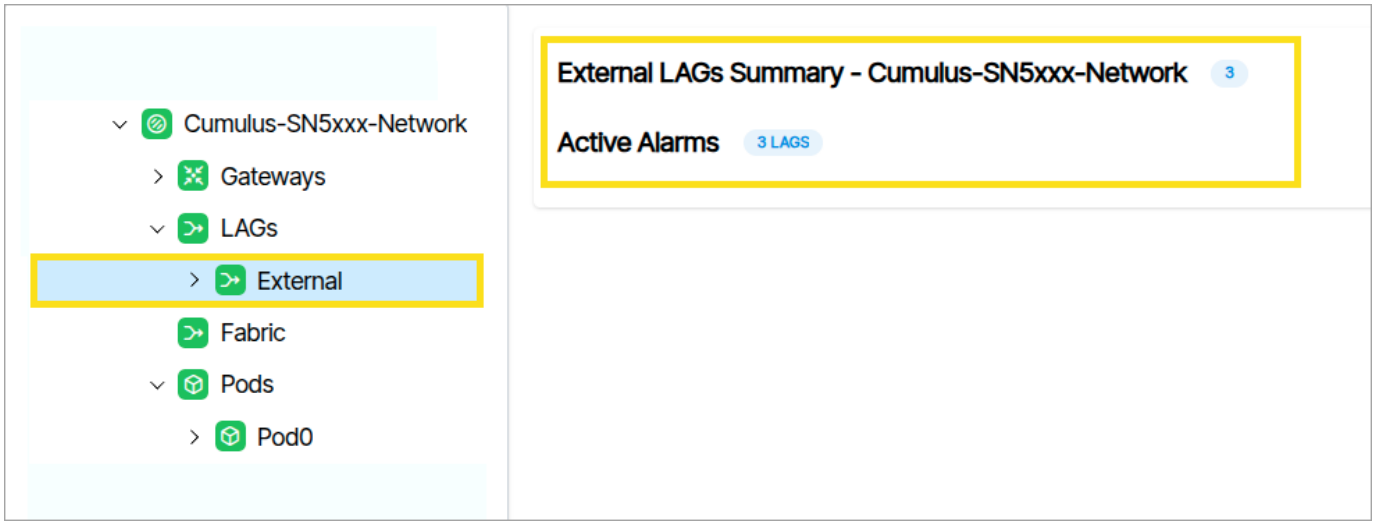
Click the Priority numbers below the DSCP numbers (image needed) to choose the range of numbers to set (image needed).

Notes

A persistent note writing section for administrators.

Container Details

Some tree view items exist only to group their children. Clicking these nodes reveals aggregate details about the child objects. The following image shows that the node named **External** contains 3 LAGs.

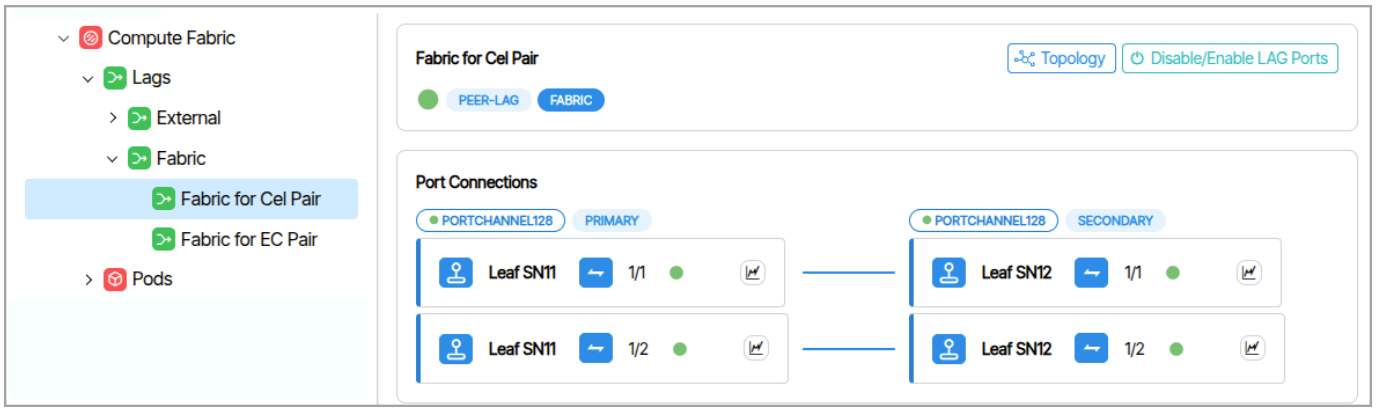


LAGs (Fabric/External)

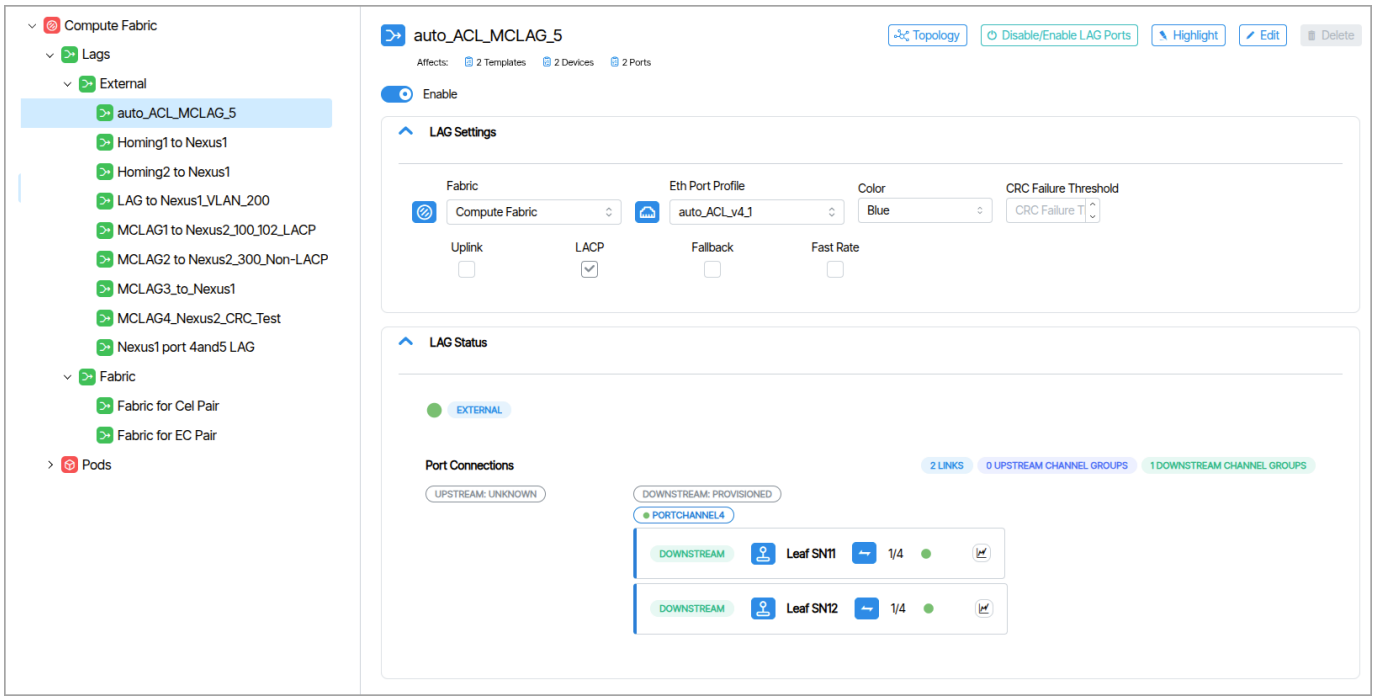
You view and manage LAG parameters from the Fabrics navigation window.

Copyright BE Networks 2026

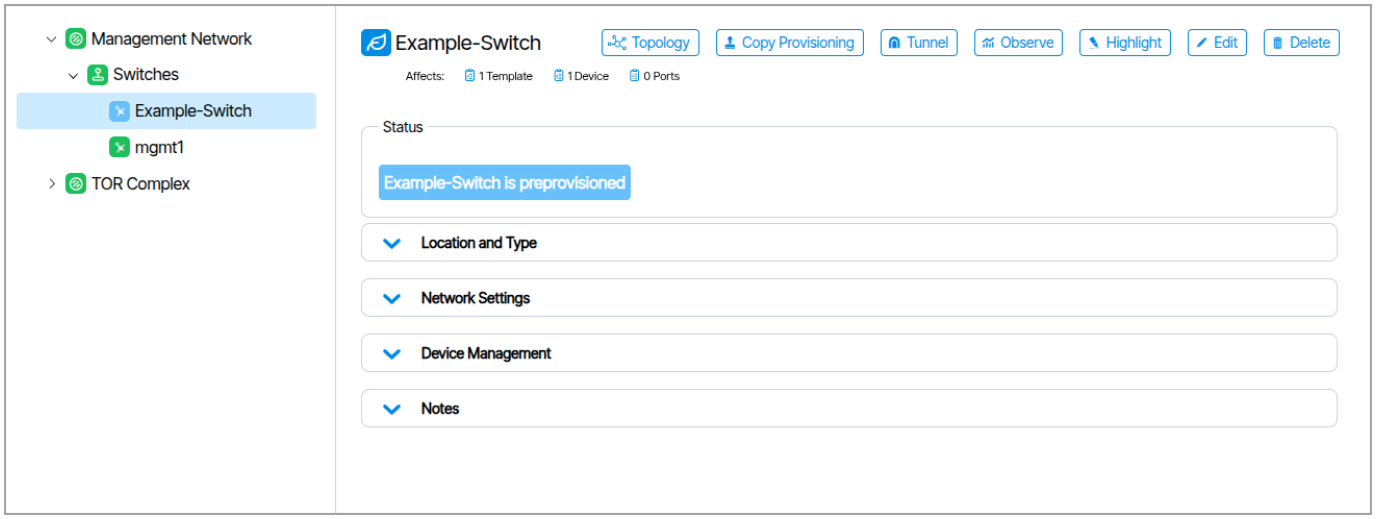
FABRIC LAG PARAMETERS



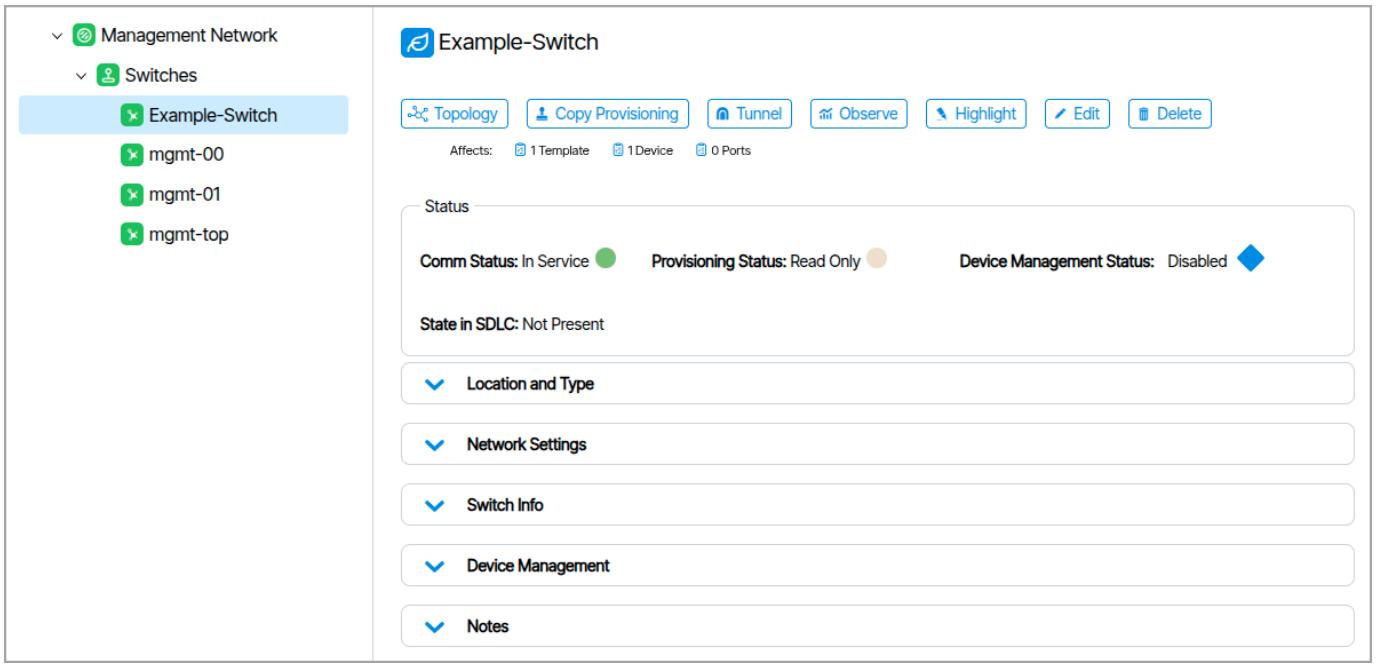
EXTERNAL LAG PARAMETERS



Switches (Spine,Leafs)



Copyright BE Networks 2026



STATUS

These indicators provide information about the device state.

- **Comm Status**: Indicates the overall communication health and connectivity state of the device within the management system.
- **Provisioning Status**: Indicates the current provisioning state.
- **Device Management Status**: When you enable Device Management parameters, this indicates the connection status.
- **State in SDLC**: Indicates SDLC behavior.

LOCATION AND TYPE

Type of device (Spine, Leaf, Superspine, Management) , POD and Rack information is listed here.

NETWORK SETTINGS

Network Settings vary depending on the device type.

SWITCH INFO

Switch Info displays device hardware details, such as make and model, and basic configuration information, including the current management IP address. This information is also viewable from Topology at the top of Switch ( ).

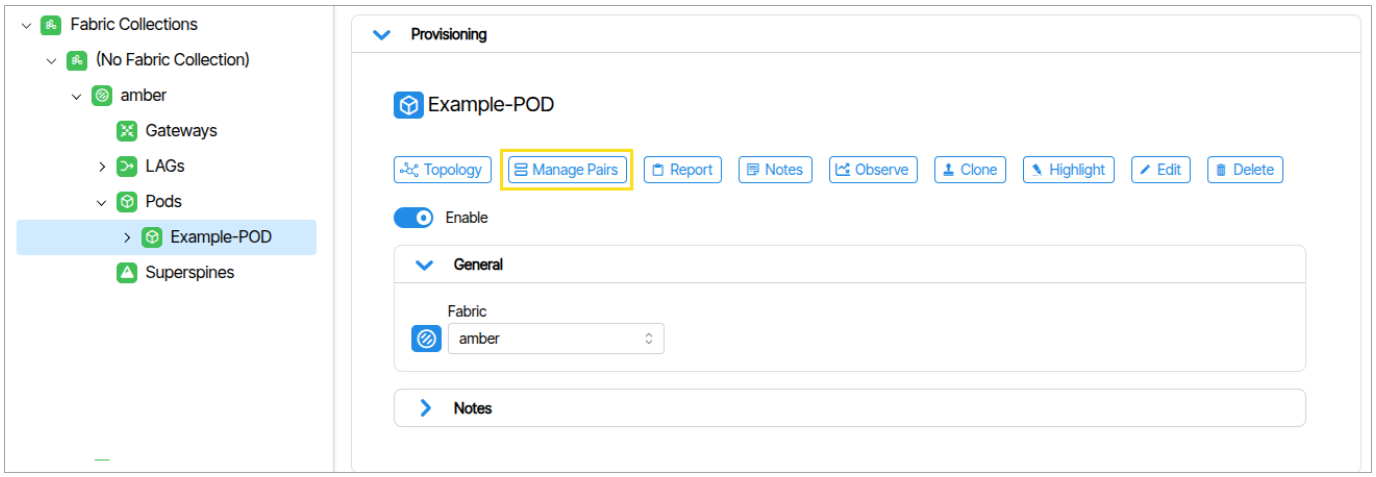
DEVICE MANAGEMENT

Device Management settings are editable and fall under Switchpoint settings. In addition to the parameters in Device Management, Switchpoint settings include all port provisioning of the switch. You can copy these values between devices using the **Copy Provisioning** feature.

How to Create a Switch Pair From Fabrics

- Determine the target devices. If the target devices do not yet exist, create them before creating a switch pair.
- Select the chosen Fabric or POD to apply the Switch Pair.
- In the upper portion of the page, click **Manage Pairs**.

Copyright BE Networks 2026



- Enter a name in the **Name** field, then select your devices under the **Switchpoint** fields.
- Use the **Multi Chassis** checkbox  to control whether the pair is a redundancy pair. The checkbox is selected by default for data center fabrics. Clear the checkbox to create a non-redundant L3 leaf switch pair.

When complete, the Fabrics tree displays the devices as a related pair with the Switch Pair icon. Peer LAG, MCLAG, and pair link status handling apply only to pairs where **Multi Chassis** is selected.



How to Create a Fabric Gateway

A **Fabric Gateway** represents each end of a planned BGP session between fabrics. For high availability, four fabric gateways are needed — one for each end of each session.

1. From **Fabrics**, select the **Fabric**.
2. Select **Gateways**.
3. In the upper right corner, click **Add Gateway**.

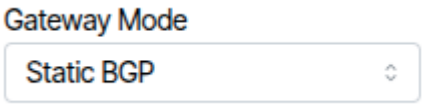
Configure each Gateway object with the following:

- The Neighbor AS number of the remote Leaf switch
- The IP address assigned to the Gateway Port Profile on that neighbor's switch
- The **Fabric Interconnect** option enabled

GATEWAY MODES🔗

Gateways can be one of several different modes. The Gateway Mode field defines how the software handles and routes traffic within the network. It determines whether routing decisions are based on a statically defined default gateway, static routes, traditional BGP routing, or new enhancements to BGP to support dynamic neighbors.

Copyright BE Networks 2026



- **Default**: A single static default route.
- **Static**: Uses manually configured static routes, no dynamic routing.
- **Static BGP**: Establishes BGP sessions with manually configured BGP neighbors. The BGP neighbors are explicitly defined by their IP addresses, and no dynamic peer discovery occurs within the subnet.
- **Dynamic BGP**: BGP neighbors are dynamically discovered within a subnet, leading to the creation of multiple BGP sessions from a single tenant gateway object.

Copyright BE Networks 2026

Leaf Spine

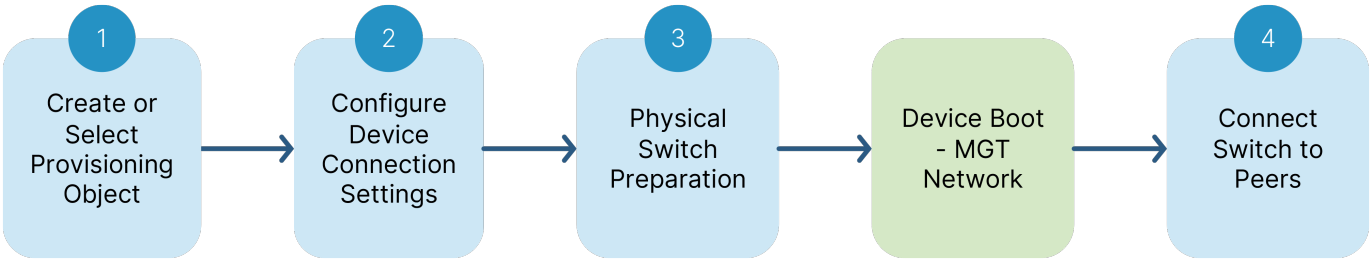
4.2.2 Device Management

This section explains how to onboard devices to a Verity system. The most efficient approach for creating the system initially is to use the bulk configuration method via a system initialization file. However, devices can also be onboarded individually through the Verity UI.

Adding Individual Devices

> **Note**
>
> This section describes how to manually add devices to the system. To add devices in bulk please see the section titled Onboarding Devices in Bulk

Adding a new spine or leaf device is a multi-step process described by the following steps:



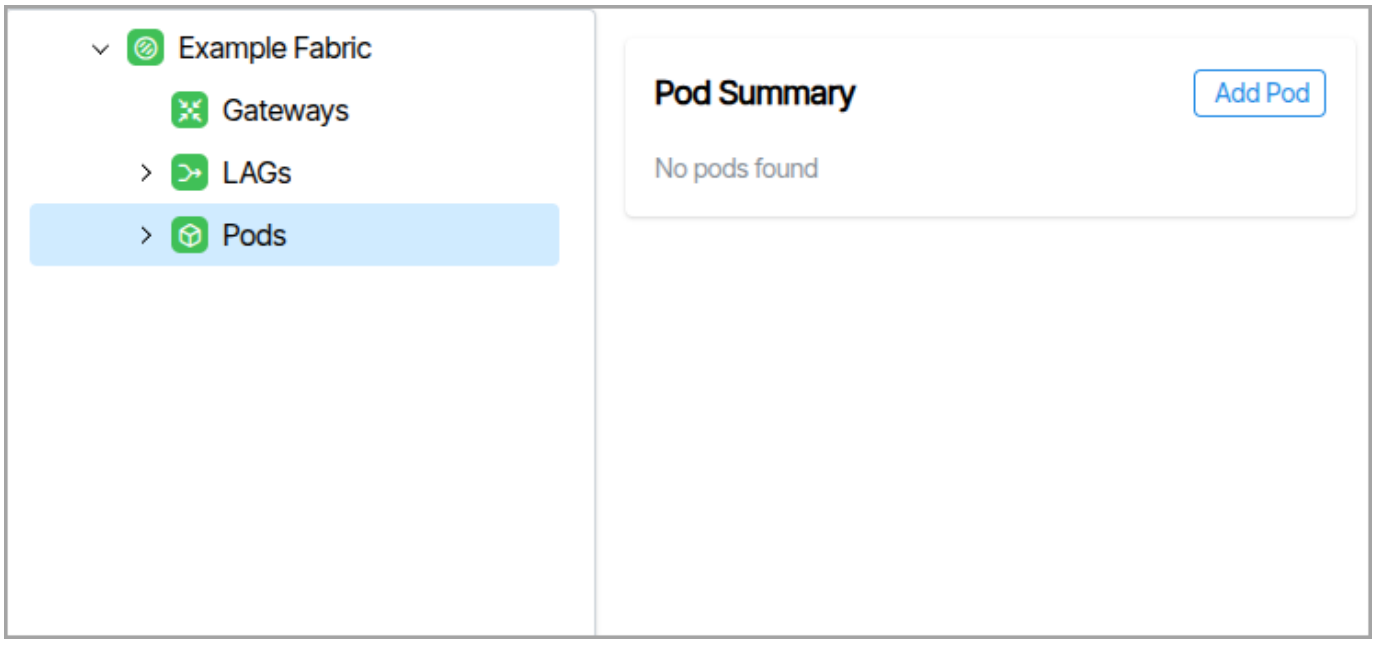
1. Create or select the provisioning object.
2. Configure the **Device Management Connection** settings.
3. Prepare the physical switch.
4. Boot the device.
5. Connect the switch to the peers.

Device Object Creation

Preprovisioning the device object enables setup of the device before the hardware arrives.

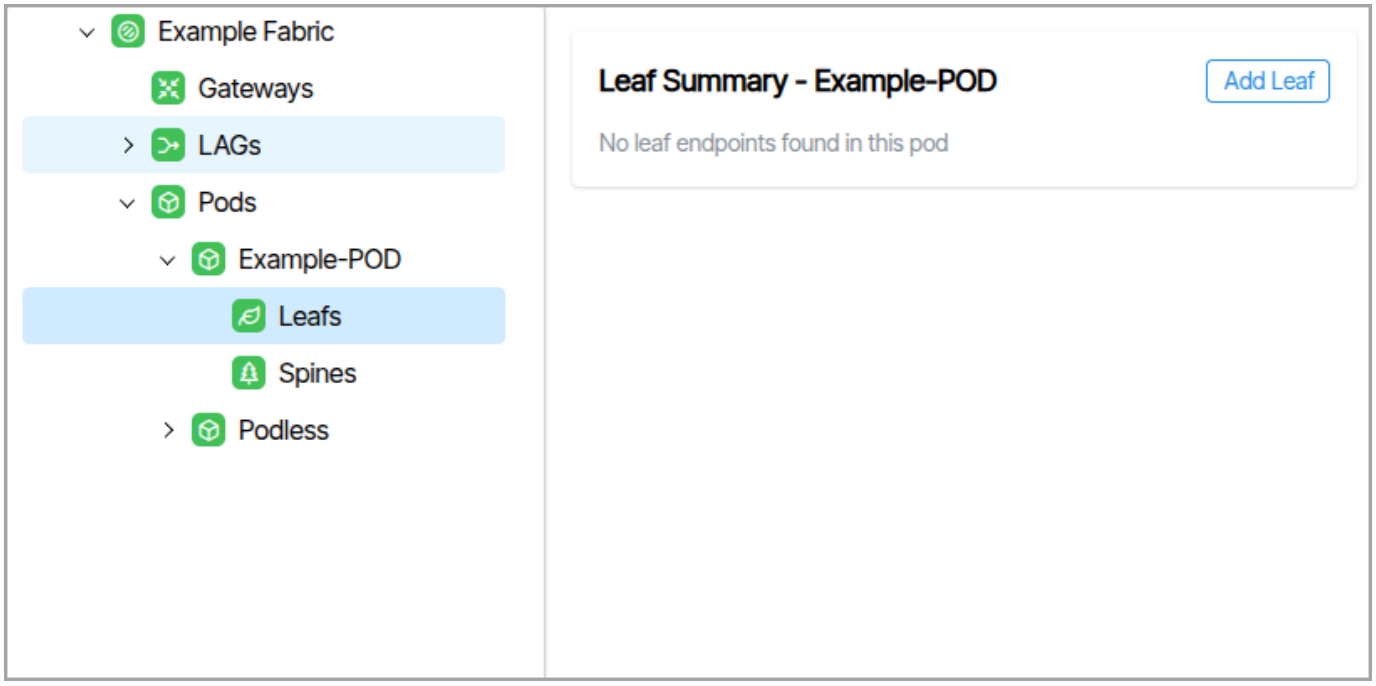
1. From the main menu go to **Fabrics** and select the fabric where you wish to add a device. If no fabric exists, *you must first create it*.
2. The type of Fabric you select determines which object types appear in the tree. In the image below, the **Example Fabric** uses the **Data Center 3 Stage Clos** type. This Fabric exposes three objects — **Gateways**, **LAGs**, and **Pods** — each represented by a collapsible tree node  .
3. The Pod object can contain Leaf and Spine devices. To create a Pod, select the **Pods** tree node, click **Add Pod**, and enter a name.

Copyright BE Networks 2026



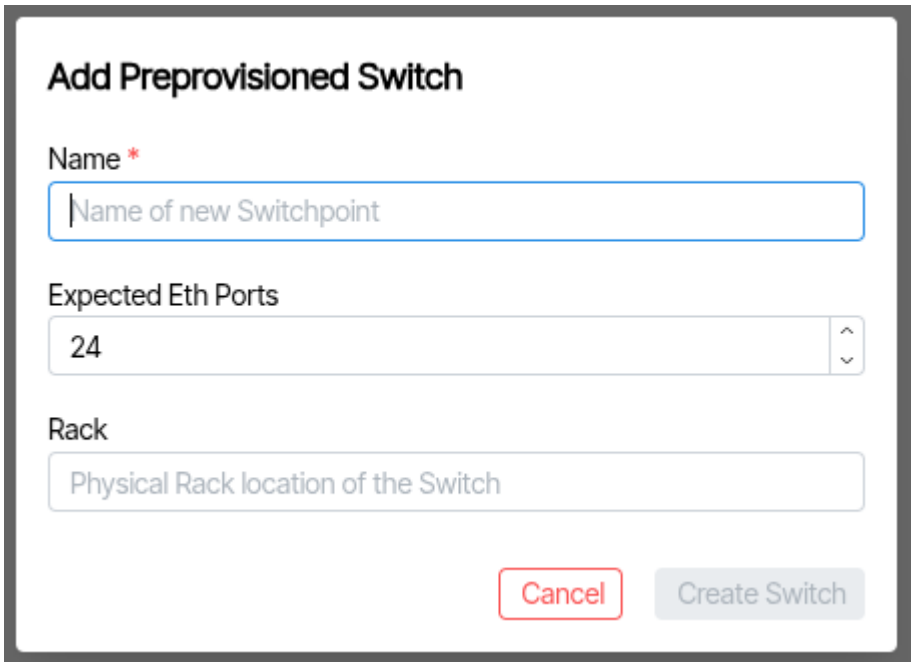
ADDING LEAF OR SPINE DEVICES

- After creating a POD, select either **Leafs** or **Spines** (depending on the device you want to create) and click the Add Leaf/Add Spine button in the upper right corner of the screen.

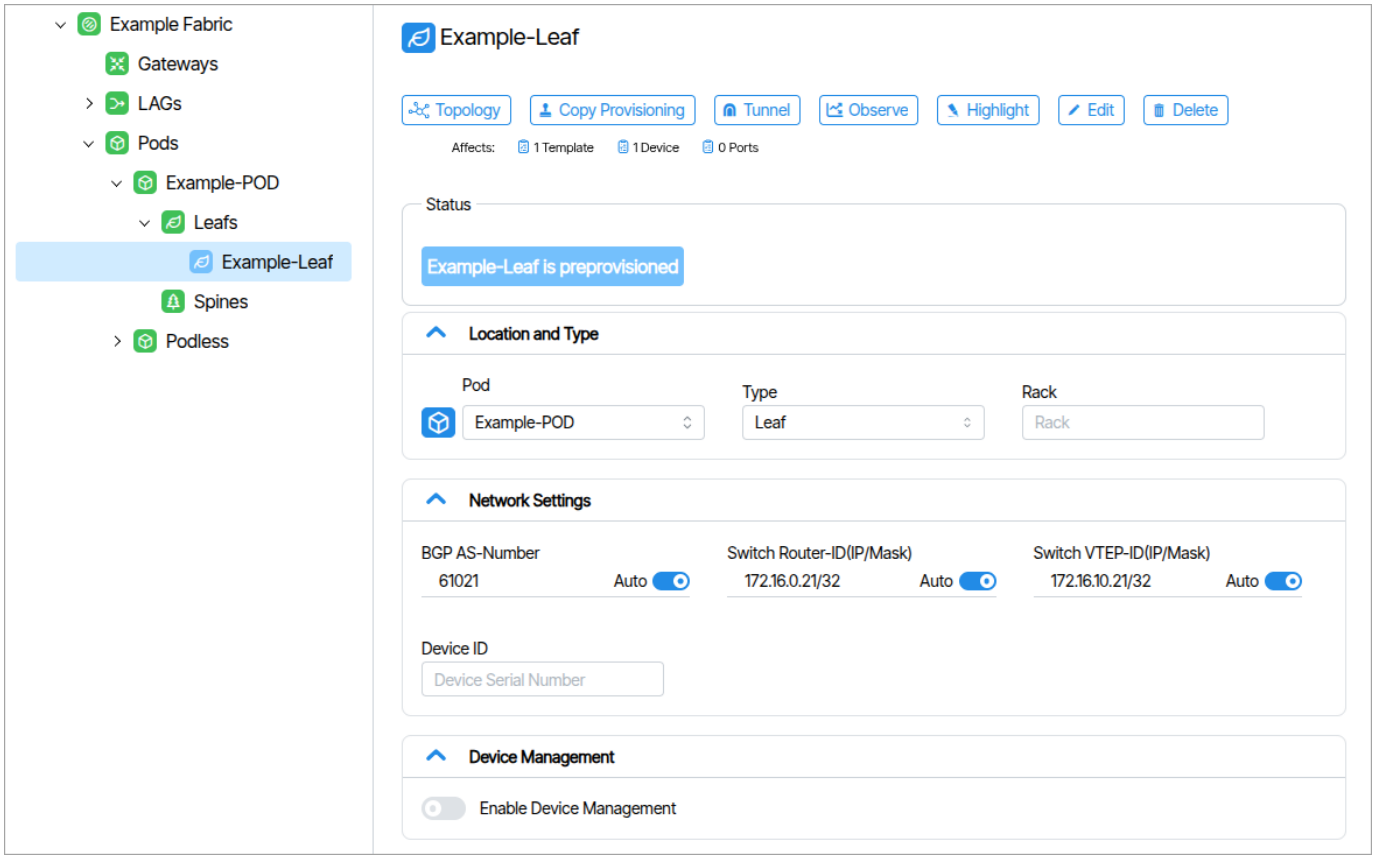


The prompt will include additional parameters to the name including the number of expected eth ports.

Copyright BE Networks 2026



The object appears blue to indicate pre-provisioned status; selecting it displays the device parameters.



Device Management

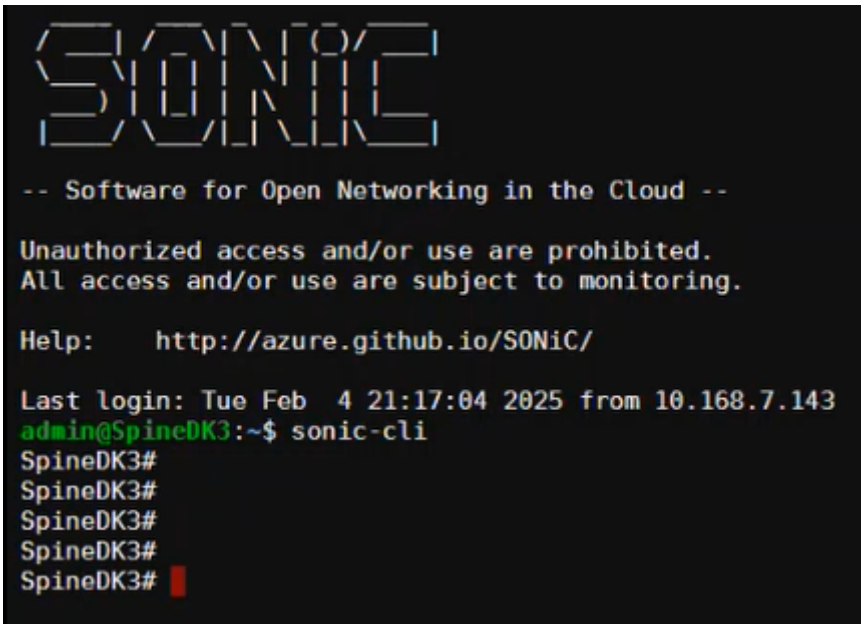
Device Management is used to create the interface between the Verity application and the switch. Based on the LLDP capabilities of the switch, it can be automatically discovered in the network topology, or it can be statically located. For systems containing many devices, the configuration settings in Device Management can be imported/exported from the Import/Export tool bench available from **Operations / Import/Export**

Copyright BE Networks 2026

CONFIGURE DEVICE MANAGEMENT SETTINGS

1. Go to **Fabrics** and select the device. Navigate to the section titled **Device Management**.
2. Enter the **LLDP Search String**. This value must be either the chassis ID or the serial number of the managed device. This value serves as a hardware identifier and is used to detect connections between managed devices. Later, this value will auto-populate in the pre-provisioned Device ID when the entries to the Device Management settings are saved.
3. Enter the **Switch IP and Mask** and **Switch Gateway**. These fields accept IPv4 or IPv6 addressing. The switch IP and gateway must use the same IP version, and when a mask is provided the gateway must be within the switch IP subnet. For data center switchpoints, the switch IP requires a mask except for HGX switchpoints. Host-only masks are not allowed, such as IPv4 `/32` or IPv6 `/128`.
4. **Enable** the Device Management settings and save the changes.
5. **Verify the Hardware Identifier**. The Device Management LLDP Search String should appear on the provisioning object in the **Device ID** field. ()

How Switch Hostnames are provisioned



The Device Management settings for each switch includes both an LLDP Search String and a ZTP Identifier field. How these fields are populated determines how the hostname is assigned to the device. The hostname can be the set to the switchpoint name or to a hardware identifier (serial number or chassis ID).

- If only the LLDP Search String is entered and the ZTP Identifier is left blank, Verity uses the LLDP Search String to create ZTP files, set the hostname and identify the device during network topology discovery.
- If both fields are filled with different values, the ZTP Identifier is used to create ZTP files, the LLDP chassis ID is used to identify the device during network topology discovery, and the hostname is set to the switchpoint name.

Switch Preparation

New Devices Starting Point

It is required that all SONiC devices being onboarded into Verity start in ONIE mode with no OS installed. This ensures that the device receives its load from the Partner Firmware Package installed on Verity and runs through the ZTP process. (If the current switch is running OS10 with DHCP enabled, OS10 is automatically removed when using DHCP options set up according to this documentation.)

Copyright BE Networks 2026

> **Reusing Devices**
>
> If devices have previously been used by other network management systems, it is essential that you completely "zero" the switch with the following commands:
>
> sudo ZTP erase -y sudo ZTP enable sudo ZTP run -y

Device ZTP

The ZTP feature is a way to remotely set up switches without having to manually configure them on a hop-by-hop basis. The process consists of two steps: ONIE Boot process and ZTP Provisioning portion.

The ONIE process starts when a switch boots up from the factory without a SONiC image and it connects to a remote server to download the necessary software image. The actual ZTP process starts with a switch that has SONiC loaded and a default configuration. The Verity system creates a series of files that are used to prepare the SONiC switch with a minimal configuration to communicate with the Device controllers within Verity. DHCP options on the Verity SDLC server tells the switches the location of the remote server (vNetC) where the required files are stored. The system supports ZTP for both the management and spine/leaf switches.

1. When ready, connect the management port of the new device to the **Out-of-Band Management** network.

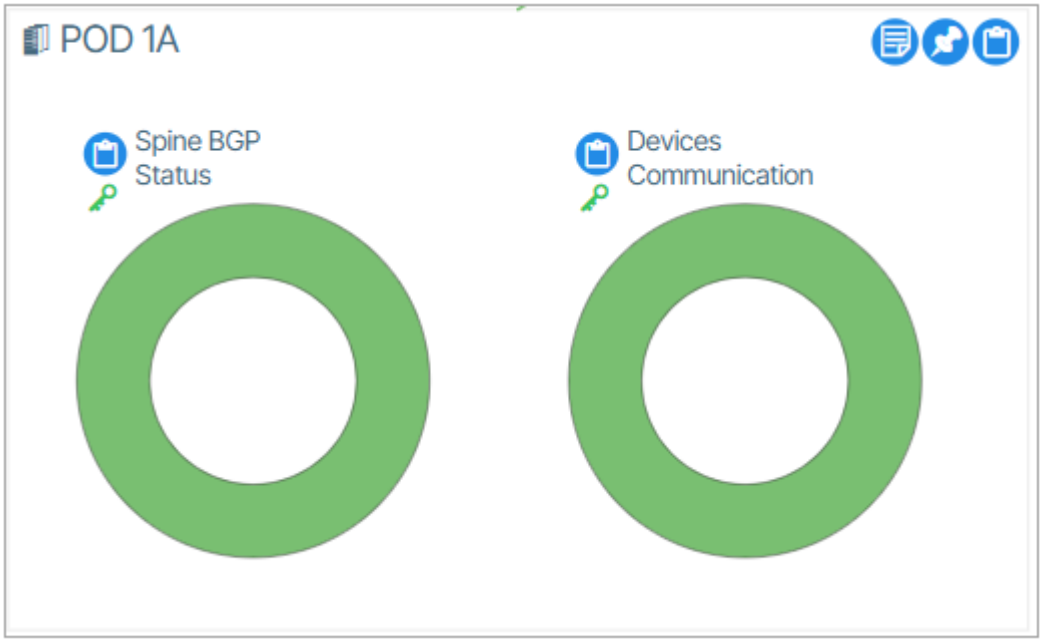
Connect Neighbors

1. To complete the addition of the new device, connect the cabling to the other relevant spine/leaf switches.
2. Go to the **Topology** view and ensure that the neighbor links become operational.

Copyright BE Networks 2026

4.2.3 PODs

In a spine-leaf architecture a POD is a grouping of leaf and spine switches. This allows for added capacity by deploying additional PODs without redesigning the entire network.



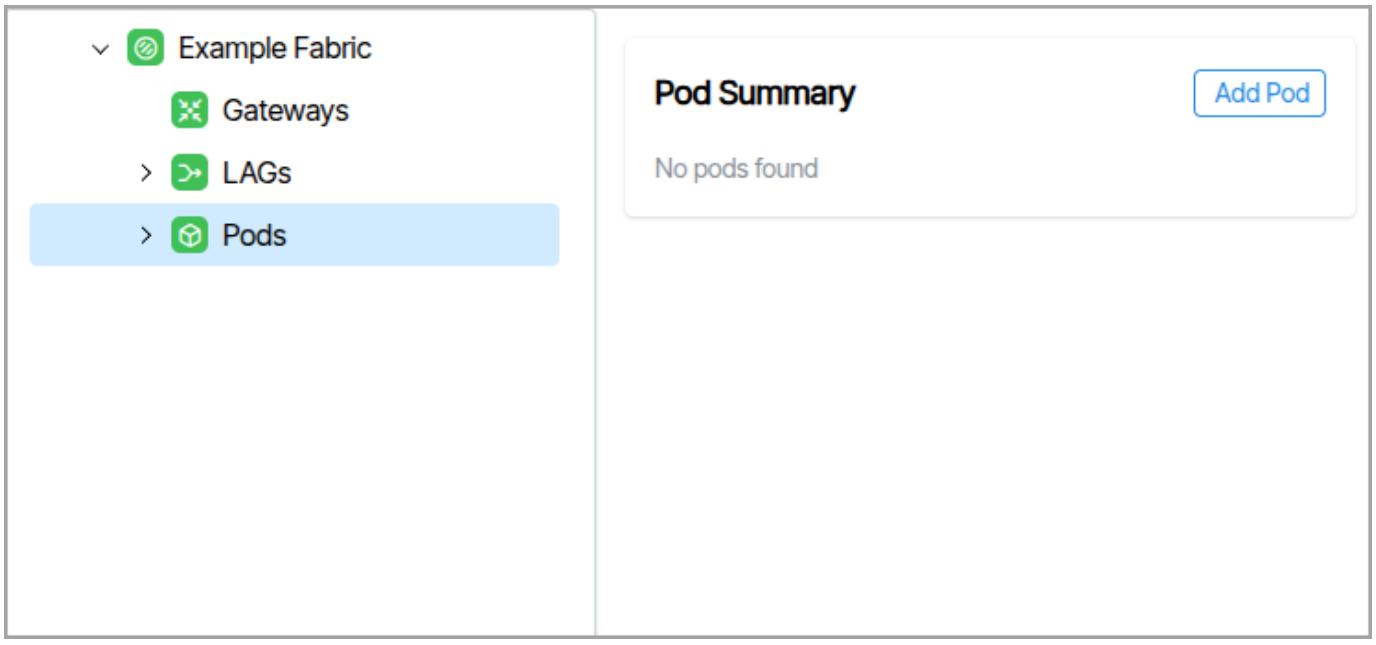
POD Reports

The **Reports** window provides a Report of all PODS via **Reports → PODs**

How to Create a POD

1. From the Main Navigation Menu select **Fabrics** and choose your Fabric.
2. Select the Pods icon.
3. In the upper right corner of the page, click **Add Pod**.
4. Type a name and click **Add Pod**.

Copyright BE Networks 2026



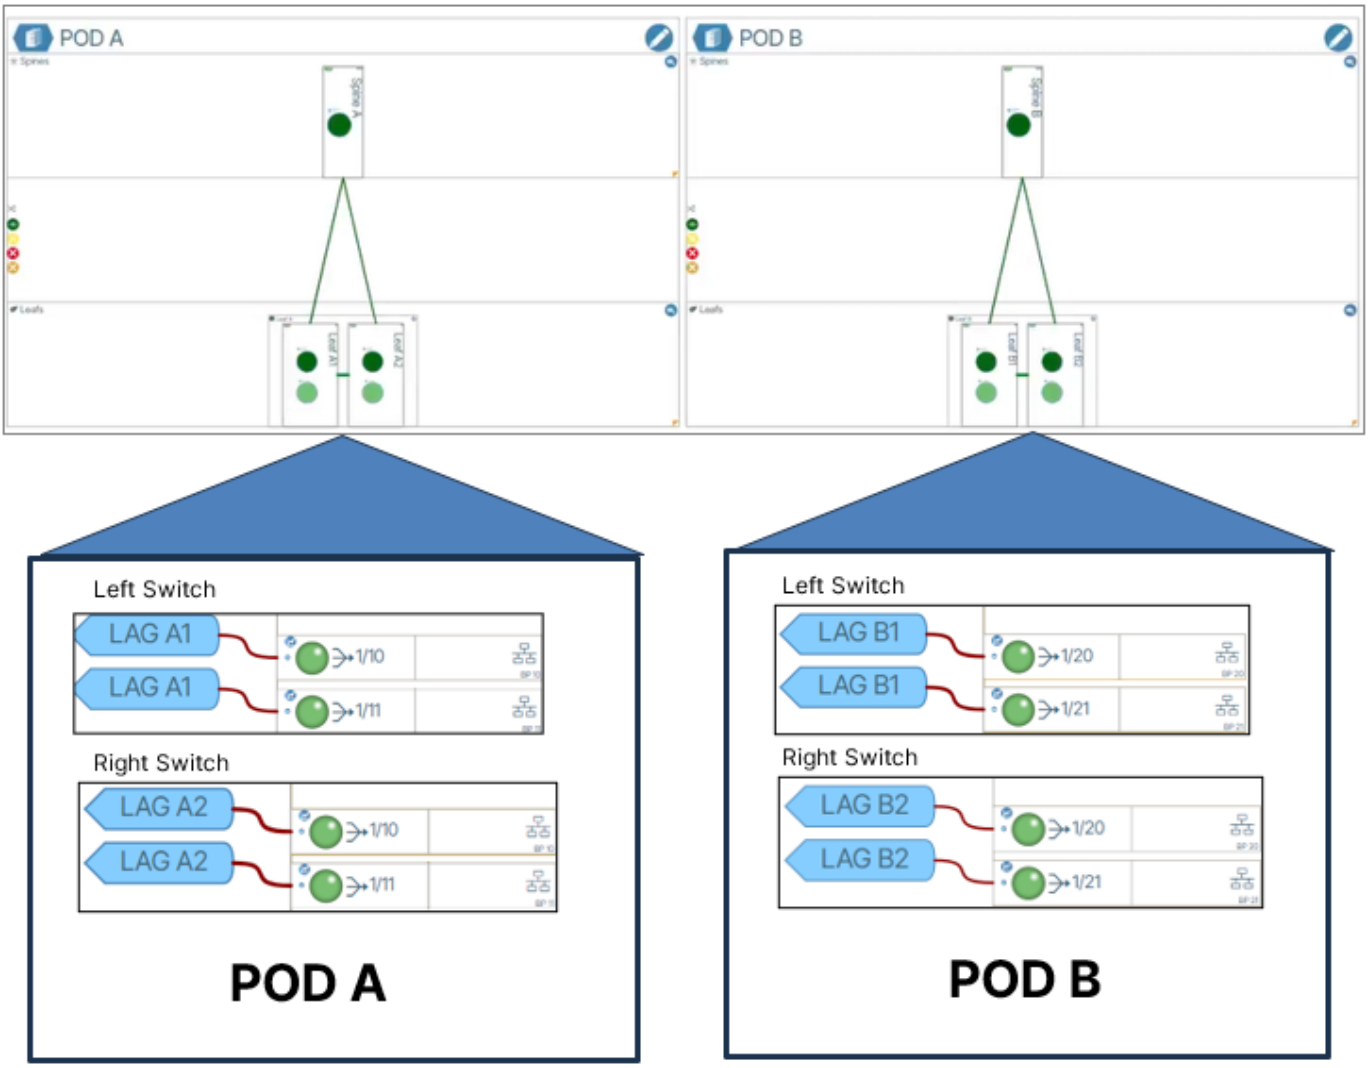
Inter-Pod Gateways

This feature enables network gateway communication between device pairs in two *separate* PODS.

To use this feature, the PODS must *not* be connected by a Super Spine.

The image below shows a basic system containing two independent PODS with each containing a device pair. The PODS are devoid of a Super Spine connecting them. It is assumed that the requisite physical connections are made (an example would be a Leaf Pair in each POD, with each leaf connected by two wires, totaling four wires and involving eight ports).

Copyright BE Networks 2026



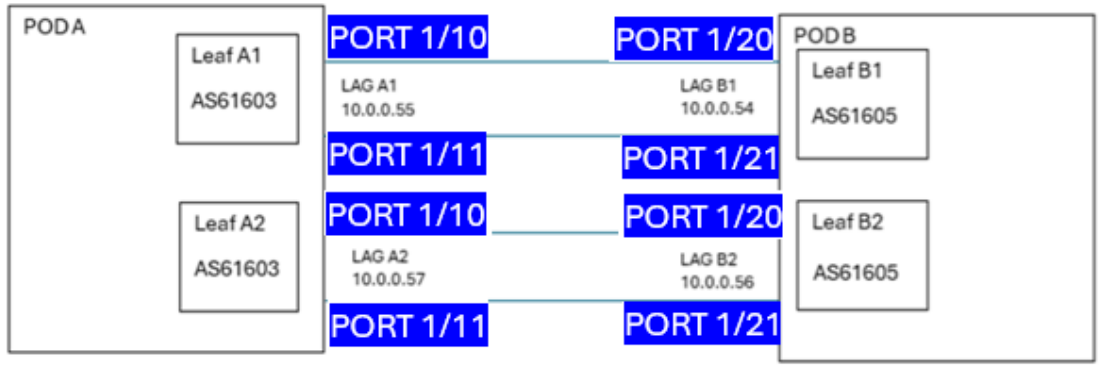
Copyright BE Networks 2026



OBJECT CONFIGURATION DESCRIPTION

The chart below presents an *example* configuration that includes 2 PODS, 4 switches and 8 switch ports. Object names, IP addresses and Autonomous System numbers will differ for your configuration.

Port numbers are presented in a blue background and are specific to the example above.



Leaf A1 – AS 61603 Leaf A2 – AS 61603	Leaf B1 – AS 61605 Leaf B2 – AS 61605
Gateway A1– Neighbor AS 61605 Neighbor IP - 10.0.0.54	Gateway B1 = Neighbor AS 61603 Neighbor IP – 10.0.0.55
Gateway A2– Neighbor AS 61605 Neighbor IP - 10.0.0.56	Gateway B2 = Neighbor AS 61603 Neighbor IP – 10.0.0.57
External LAG A1 Applied to two Ports on Leaf A1	External LAG B1 applied to two ports on Leaf B1
External LAG A2 Applied to two Ports on Leaf A2	External LAG B1 applied to two ports on Leaf B2
Gateway Profile A1 Applied to LAG A1 – Source IP 10.0.0.55/31	Gateway Profile B1 Applied to LAG B1 – Source IP 10.0.0.54/31
Gateway Profile A2 Applied to LAG A2 – Source IP 10.0.0.57/31	Gateway Profile B2 Applied to LAG B2 – Source IP 10.0.0.56/31

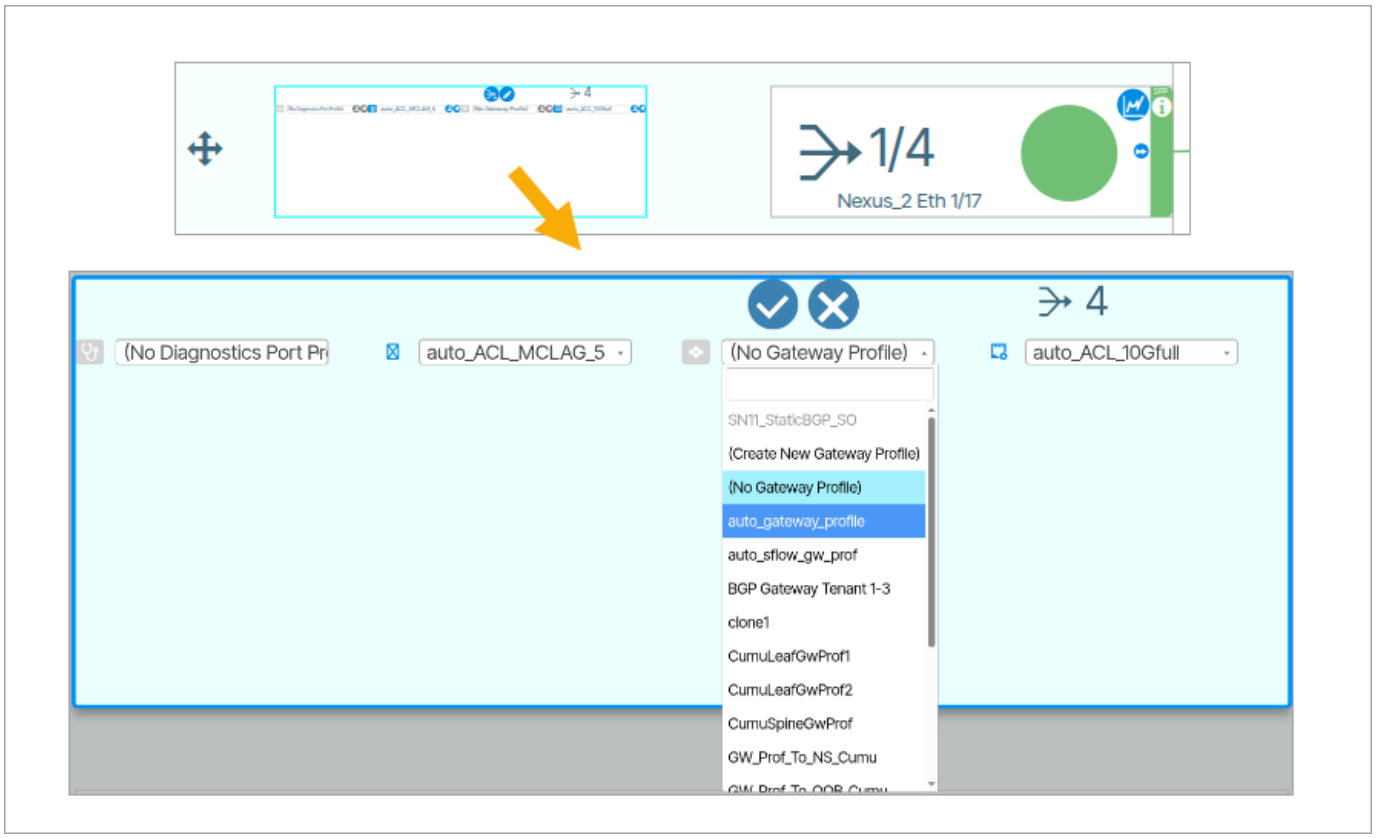
Copyright BE Networks 2026

CREATE A GATEWAY PROFILE

1. Go to **Templates / Gateway Profiles**.
2. Create the needed Gateway Profiles. (layer-3-templates.md/#gateway-profiles)

ADDING A GATEWAY PROFILE TO A LAG

Adding a gateway profile to a LAG is needed to create the connection. To add a Gateway Profile to a LAG, set the Gateway Profile in the **Gateway Profile for Eth** field on the device port.

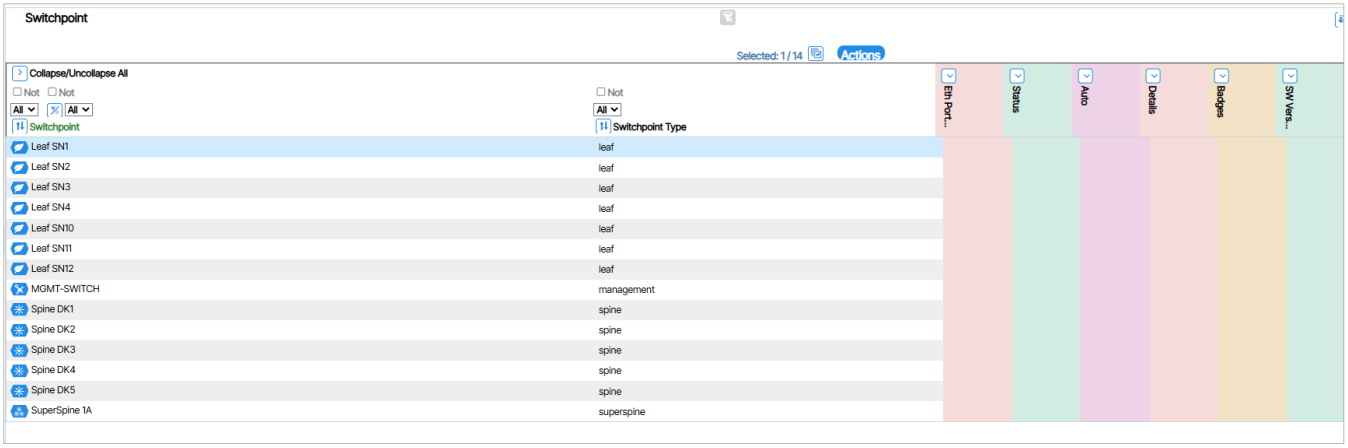


Copyright BE Networks 2026

4.2.4 Switchpoint

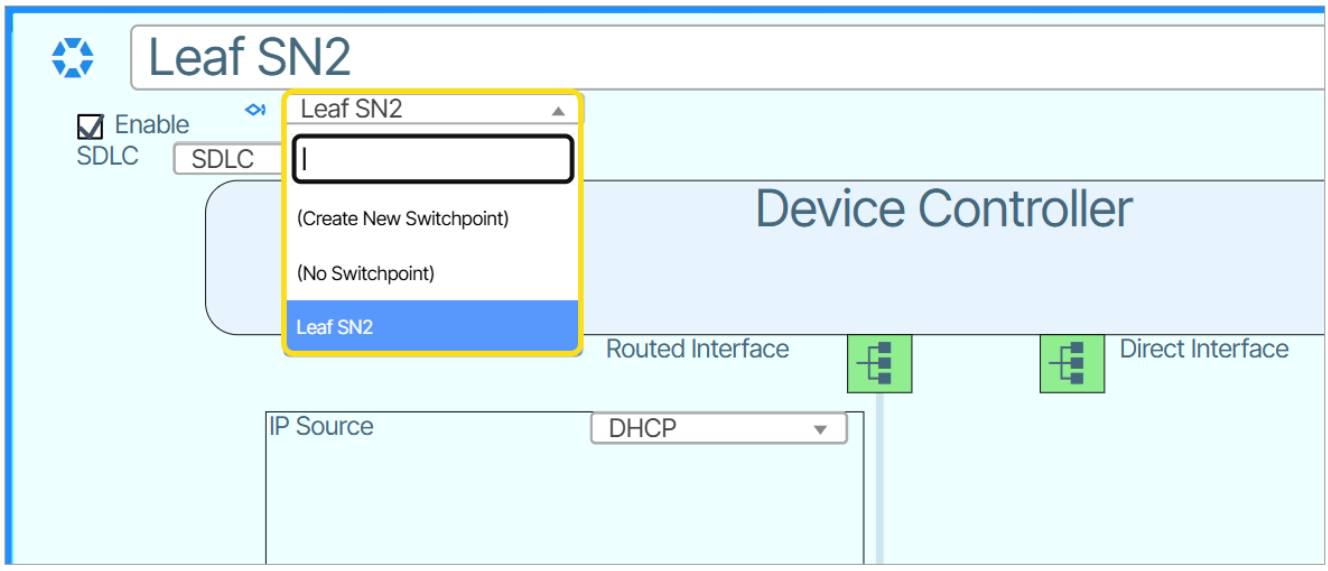
A **Switchpoint** represents a logical location for a device in the network and defines a connection point—such as where leaf 1 connects to spine 3. By designing a network using Switchpoints, operators keep the logical design independent of specific devices, allowing them to insert a device into the switchpoint by referring to either its serial number or a chassis id.

The **Device → Switchpoint** window provides a Report of all **Switchpoints**.



How to Assign a Switchpoint

1. Select the **Switchpoint** field and assign the Switchpoint.



Copyright BE Networks 2026

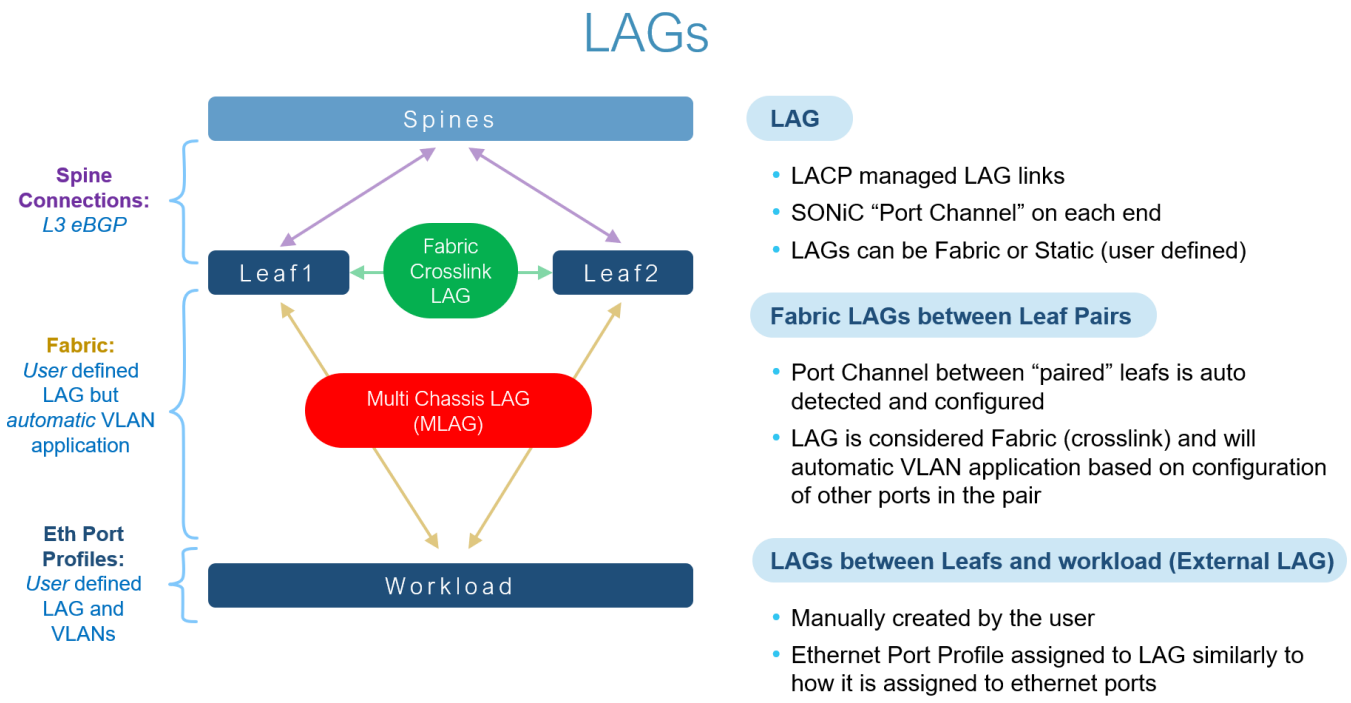
ESI Fabric LAG MLAG Multihoming

4.2.5 Link Aggregation (LAGs)

Link aggregation is a popular method for grouping multiple ethernet connections into a single logical link. Traffic across this link is automatically load balanced across the grouped (bonded) links. LAGs can extend from single leaf switches or can be configured across a pair of leaf switches. When a leaf pair is created and LAGs are assigned to ports on each switch, the LAG is considered an "MLAG". MLAG, or Multi-chassis Link Aggregation, is a networking technology that aggregates multiple physical ports across two or more switches into a single logical link, thereby providing increased bandwidth, redundancy, and resilience. This allows servers and other network devices to connect to different switches and appear as if they are connected to a single switch, improving network performance and fault tolerance.

Automation

The interconnect between Leaf Pairs (Fabric Crosslink LAG shown in diagram below) is automatically recognized and configured by Verity. This section explains how to create external LAG connections to the workload.

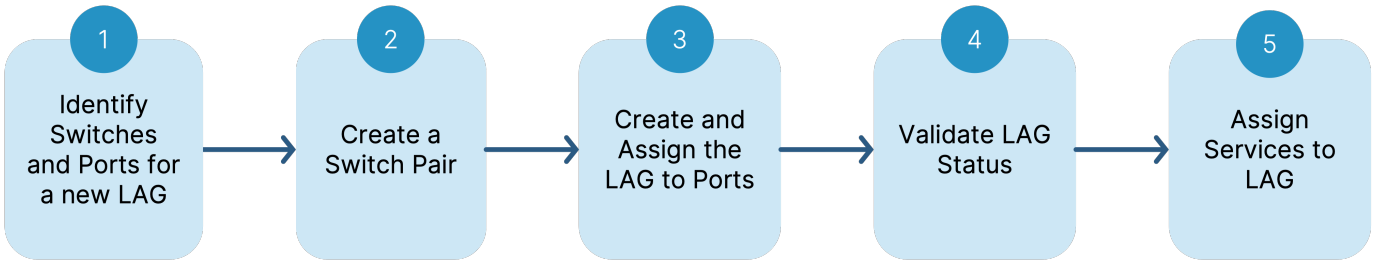


4.2.6 Link Aggregation Process

Although there are several ways of completing the process, an easy to understand method requires these steps:

1. Identify the composition of the LAG, including the leaf pairs and ports you intend to use in its configuration. (Assuming you are implementing MLAG)
2. Create the external LAG object, name it, and assign it to the port(s).
3. Select the Tenant services and create an Eth Port profile and assign it to the external LAG.

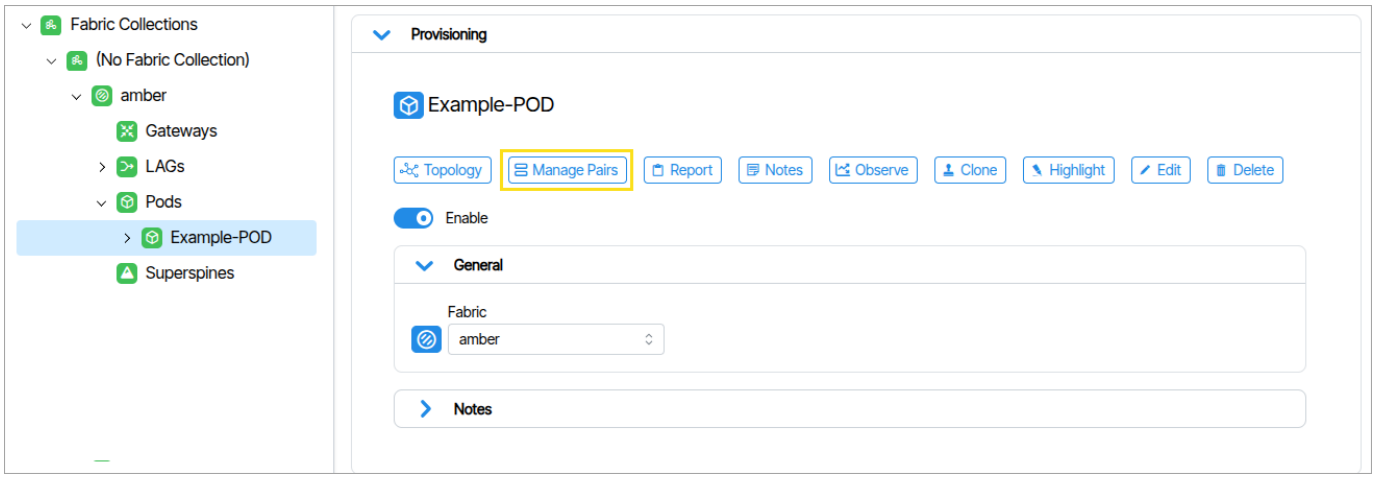
Copyright BE Networks 2026



The following example below shows how to create an external LAG relationship between two leaf ports. These actions can be used to establish different LAG configurations depending on your needs, such as aggregating multiple ports on a single leaf switch to boost endpoint bandwidth.

How to Create a Switch Pair From Fabrics

- Determine the target devices. If the target devices do not yet exist, create them before creating a switch pair.
- Select the chosen Fabric or POD to apply the Switch Pair.
- In the upper portion of the page, click **Manage Pairs**.



- Enter a name in the **Name** field, then select your devices under the **Switchpoint** fields.
- Use the **Multi Chassis** checkbox  to control whether the pair is a redundancy pair. The checkbox is selected by default for data center fabrics. Clear the checkbox to create a non-redundant L3 leaf switch pair.

When complete, the Fabrics tree displays the devices as a related pair with the Switch Pair icon. Peer LAG, MCLAG, and pair link status handling apply only to pairs where **Multi Chassis** is selected.



Copyright BE Networks 2026

How to Create a LAG From Fabrics

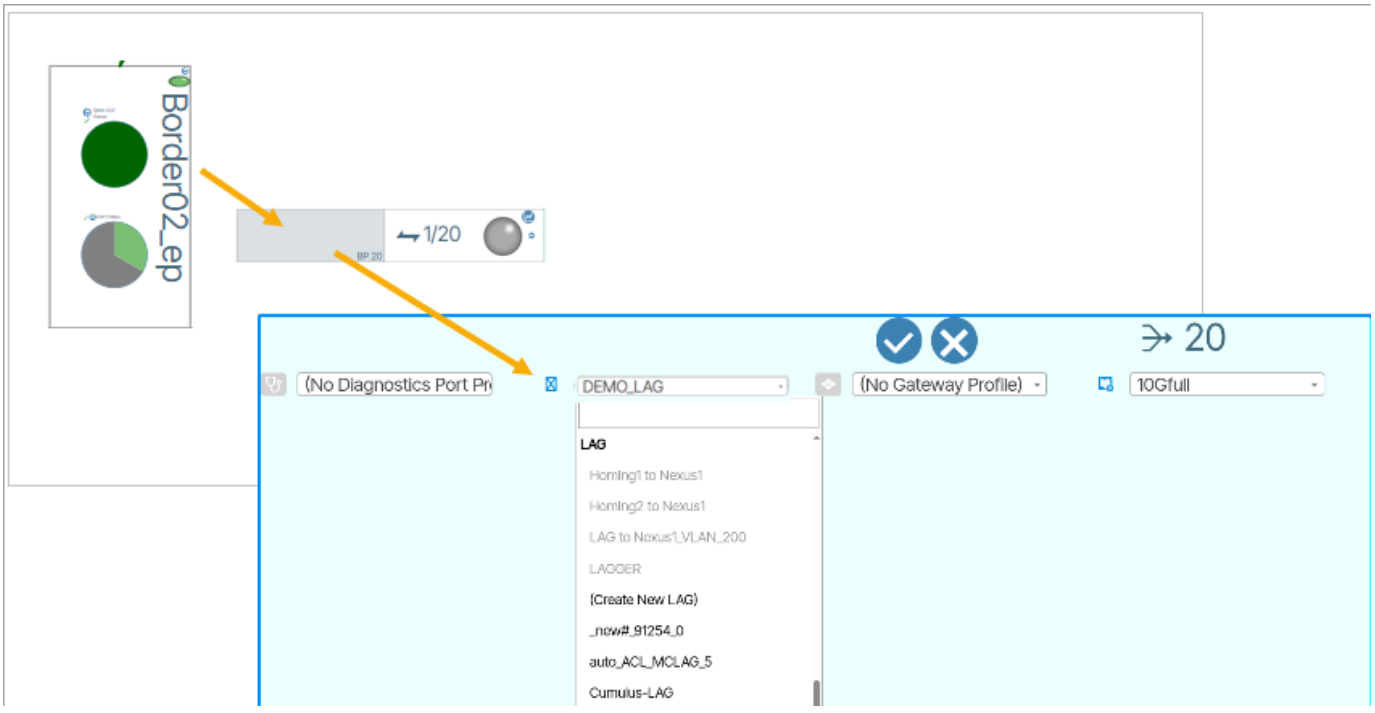
LAGs are categorized as either External or Fabric. The system auto-detects Fabric LAGs, while External LAGs can be added manually.

- From the **Main Navigation Menu** select **Fabrics**.
- Expand your Fabric, and go to LAGs
- Click External, and then click Add LAG button
- Give it a Name and Click Add LAG
- Select your new LAG
- Update/Verify the settings, set an Eth-Port Profile for the LAG
- Save the Settings.
- Apply the LAG to the interfaces on the Leaf switch you want to use this LAG.

Assign the LAG to Port

1. Select a port on your chosen device and zoom in on it. Open its parameters by clicking the space to the left of the dividing line.
2. Click the **Edit** ( ) button. On the **Eth-Port-Profile or LAG for Eth Port** drop-down menu, select the LAG ( ).
3. Click the **Save** button ( ) to save your work.

Assign LAG to Second Switch Port



1. On the second chosen switch in the pair select the desired port and open the Provision Port settings. Enable editing by clicking the editing button.
2. In the **Eth-Port-Profile or LAG for Eth Port** form field set the name of the LAG you previously created.
3. Click the **Save** button to save your work.

Validate LAG Status Display

1. As the system creates the channel groups and communicates over LACP, there will be a delay in the system for up to 5 minutes, after which the status display will turn green .

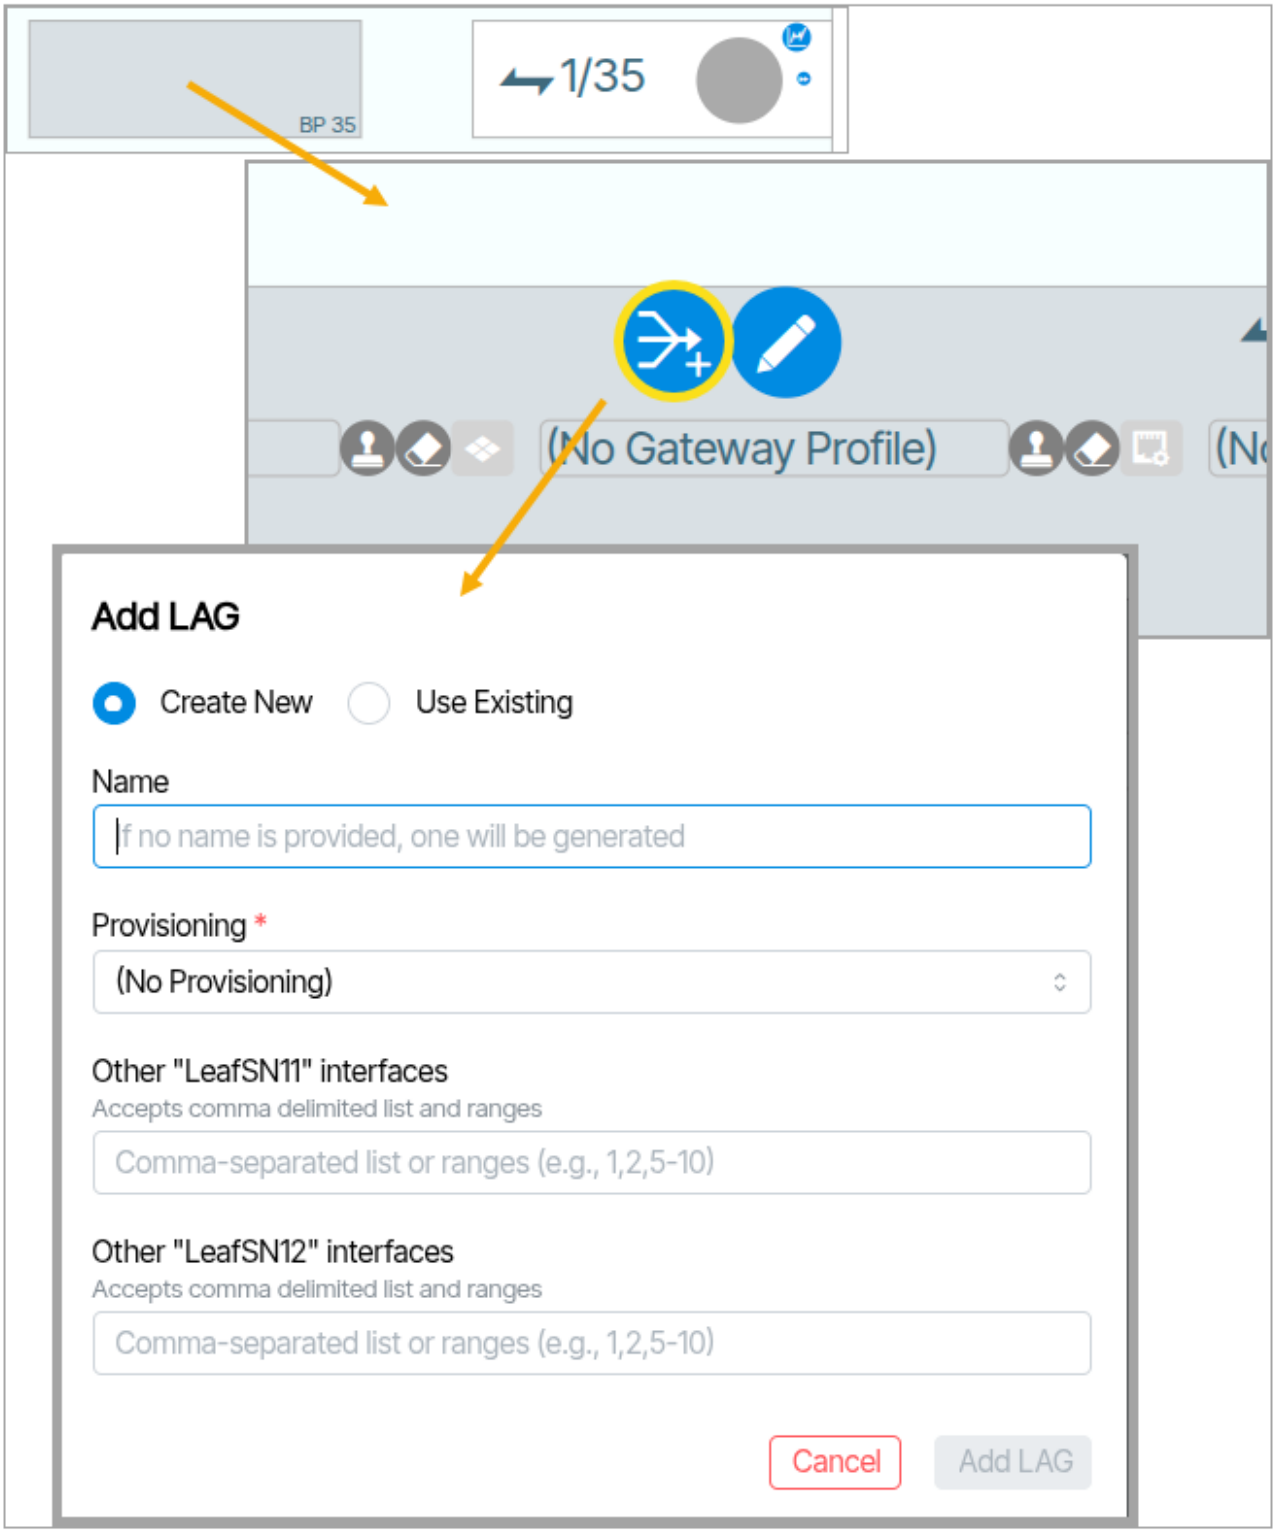
Copyright BE Networks 2026

2. If you encounter problems and the icon remains red or yellow, hovering over the status icon provides more information.

ASSIGN SERVICES TO THE LAG VIA “ADD LAG” BUTTON

1. Zoom into the settings of a port you want to assign the LAG to .
2. Click the **Add LAG** button.
3. Fill out the form with the following LAG settings:

Copyright BE Networks 2026



Copyright BE Networks 2026

Setting	Description
Create New / Use Existing	Determines if you are creating a LAG or using an existing LAG. If you choose **Use Existing,** the form field presents you with a list of existing LAGs to choose from.
Name	Sets the name of the LAG.
Provisioning	The Eth Port Profile (the service(s)) you assign to the LAG.
Other 1	The port number(s) (separated by commas if more than one) of the listed switch to assign the LAG to.
Other 2	The port number(s) (separated by commas if more than one) of the listed switch to assign the LAG to.

PROVISION CHANNEL GROUP ENDPOINTS AND SET PHYSICAL CONNECTIONS

The next step is to provision channel endpoints. This step is outside the scope of these instructions. The last step in the process is to physically connect the devices.

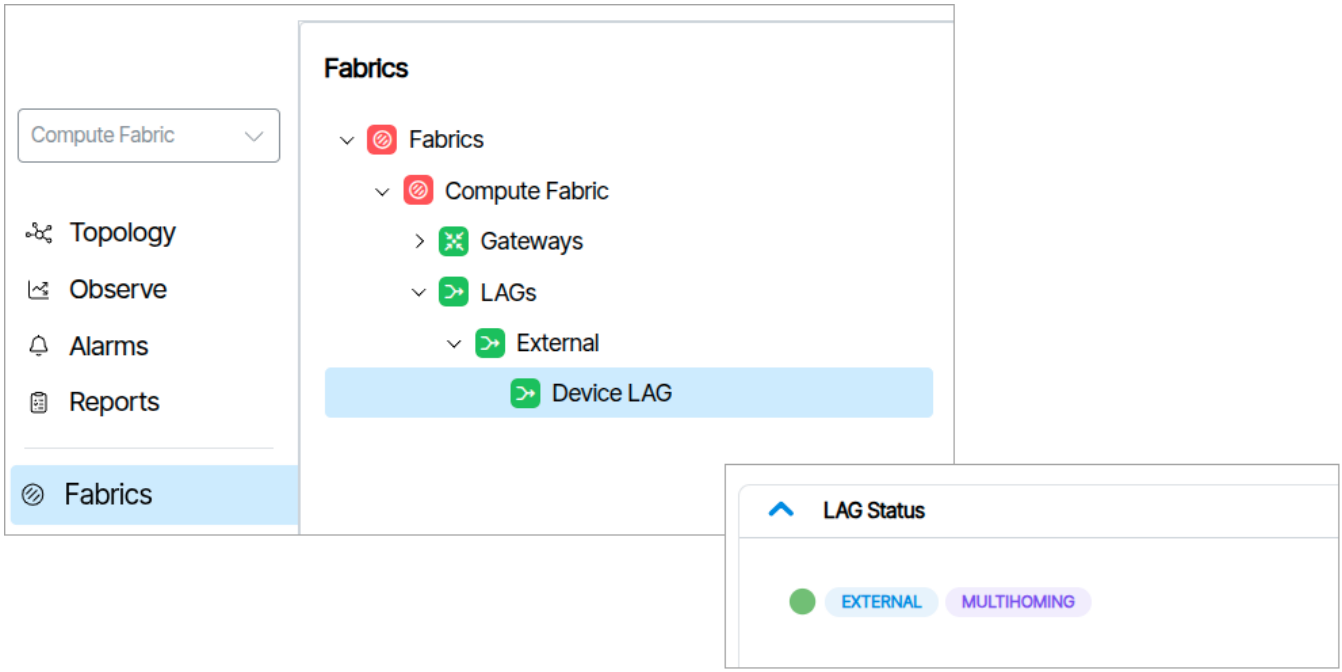
4.2.7 ESI Multihoming

A **ESI Multihoming** is a networking technology that allows multiple switches (typically leaf switches in a data center network) to form a single logical link aggregation group (LAG). Its configuration closely resembles that of a switch-pair MLAG, with the key distinction being the inclusion of up to *four* switches, significantly enhancing network redundancy and scalability. The number of ports that can be assigned a Multihoming ESI group is indeterminate. The number of switch devices is limited to four.

ESI vs MLAG

LAGS assigned to two or more *non-paired* Leaf Switches use ESI Multihoming protocol instead of proprietary MLAG implementations. Multihoming is a feature that is automatically enabled, requiring no setup from the customer.

If a LAG is using ESI Multihoming, it will display text that says **Multihoming**.



Copyright BE Networks 2026

Relationship to Leaf Pairs

The multihoming feature is implicitly disabled on all leaf switches configured as leaf-pairs. Any leaf with multihoming in use cannot be made part of a leaf pair until the multihoming configuration is removed. This prevents any conflict between the MLAG functionality on leaf pairs and the multihoming functionality.

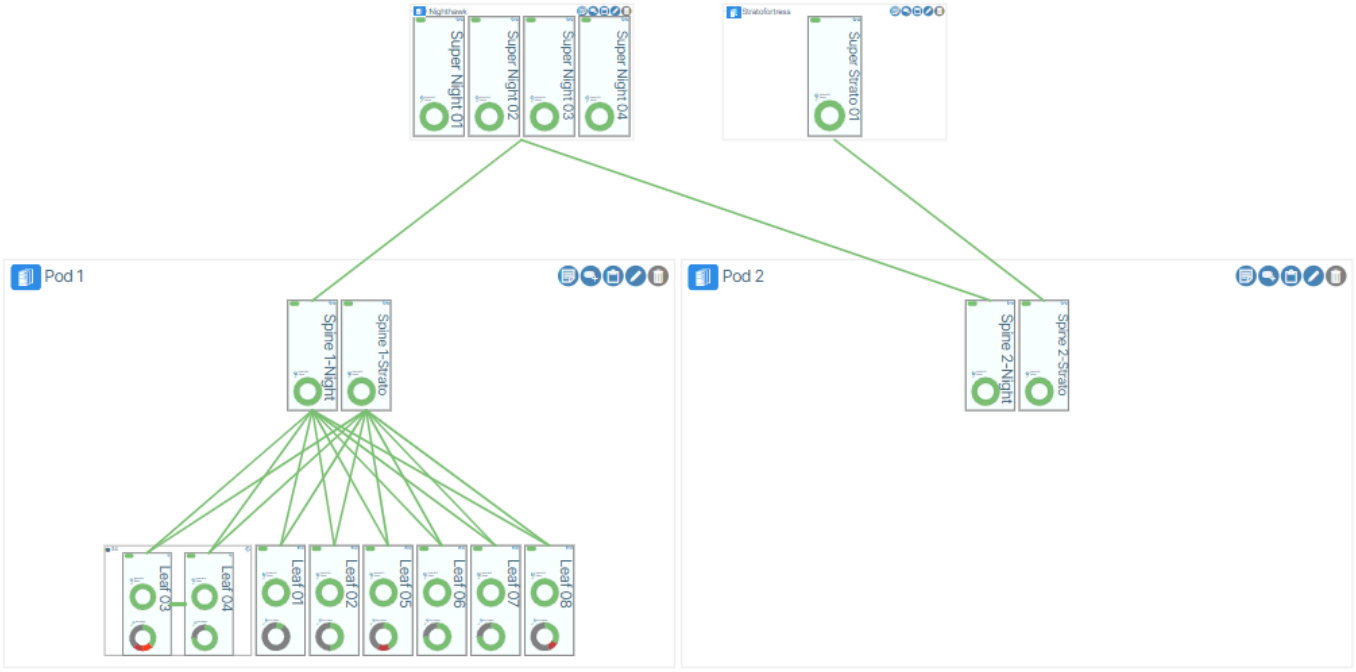
CONFLICT RESOLUTION

When attempting to pair a Multihoming switch, a conflict arises, and an error message is displayed that gives you the choice to disconnect the Multihoming connection(s).

Copyright BE Networks 2026

4.2.8 Spine Plane Architecture

Spine plane architecture enables large data centers to grow by dramatically increasing available switch ports. This form of architecture is built specifically for heavy server-to-server traffic patterns found in cloud computing, AI/ML applications, and virtualized systems where most data flows horizontally between servers rather than in and out of the data center.



Spine Plane Integration

In Verity, switch devices are organized according to their role in the system. Spine and leaf devices are grouped into collections called Pods, while Superspines are grouped into collections called Spine Planes. A Spine Plane integrates into the network architecture by connecting its containing Superspines to a spine switch within a POD, forming a dedicated inter-layer fabric. Each POD contains a single spine switch that establishes physical links to all Superspines in its assigned Spine Plane. This design creates a clear hierarchy, with the Spine Plane layer positioned above the POD layer to enable scalable, centralized interconnection between PODs.

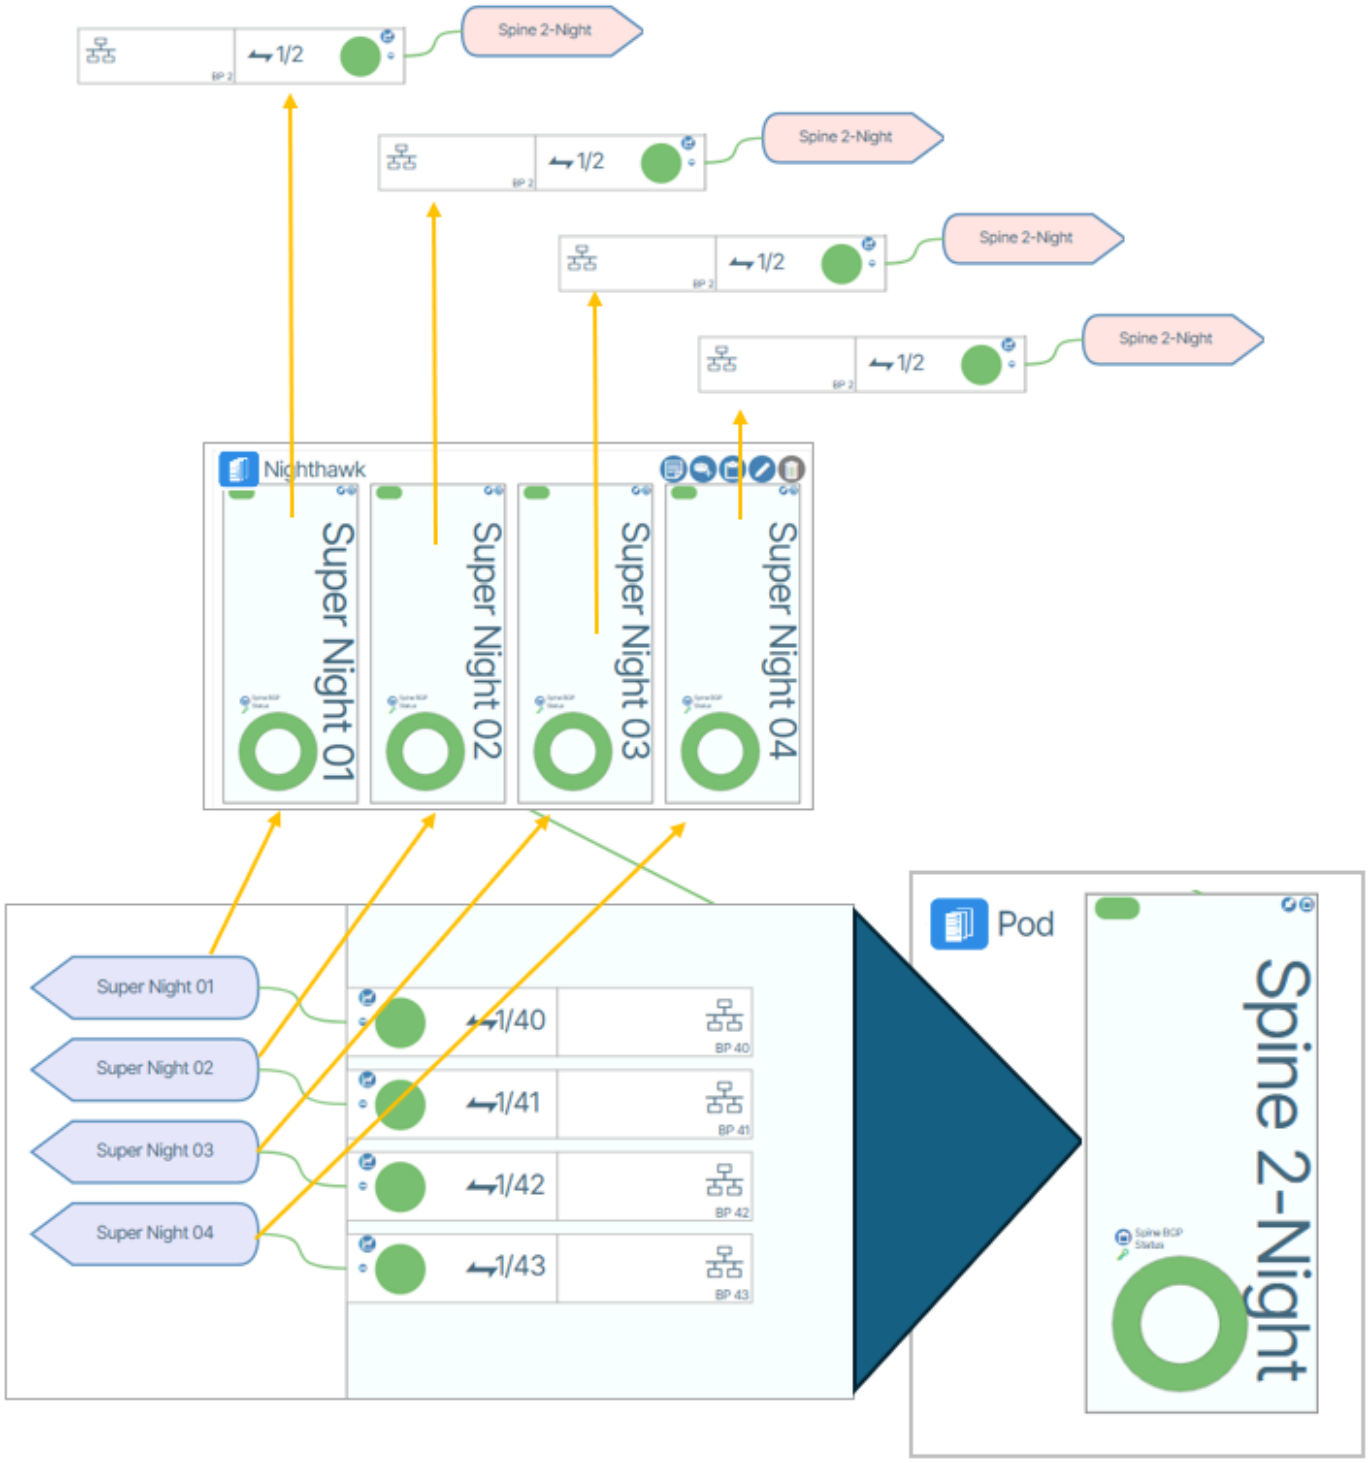
SUMMARY

- PODs are composed of Spine and Leaf switches
- Spine Planes are composed of Superspines
- One Spine per POD connects upward to all Superspines in its assigned Spine Plane
- Multiple PODs can connect to the same Spine Plane, enabling inter-POD communication

Creating a Spine Plane

Before creating a Spine Plane, the Superspine(s) and Spine hardware must be physically connected. The image below is a graphical representation of these physical connections.

Copyright BE Networks 2026



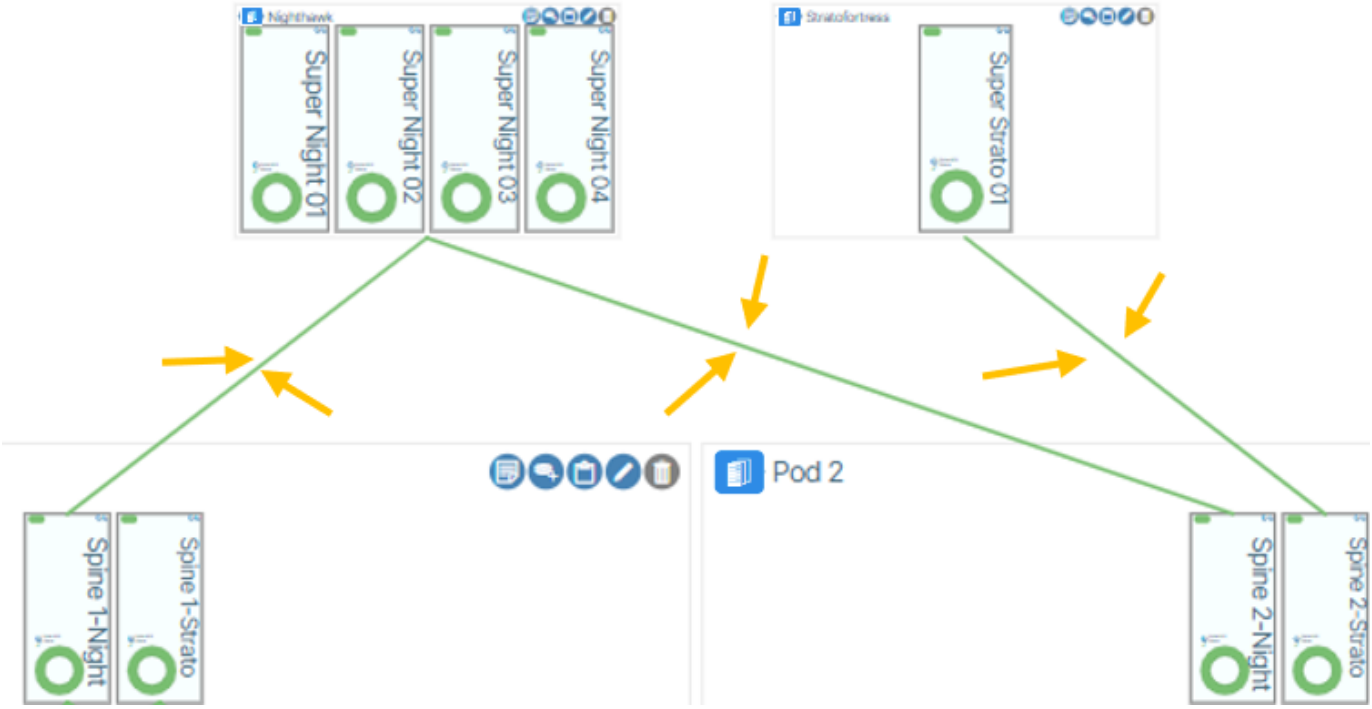
Create a Spine Plane from Fabrics

1. Click **Fabrics** in the main navigation panel.
2. Select **Fabrics** in the fabrics tree, then click **Add Fabric**.
3. Select **Multiplane** from the **Fabric Type** menu.
4. Enter a name for the fabric.
5. Configure the **Multiplane** underlay settings.
6. Select the newly created fabric.

Copyright BE Networks 2026

7. Under **Features**, enable **Allow All Underlay Connections** to permit Spine-Spine connections.
8. Under **Spine Planes**, click **Add Spine Plane**.
9. Enter a name for the **Spine Plane**.
10. Under the new Spine Plane, click **Add Superspine** to add a Superspine.

When the configuration is complete, Verity will display virtual cables between the Spine Plane(s) and Spine(s) representing the relationship between the devices.



Copyright BE Networks 2026

4.2.9 Traffic Mirrors

Traffic Mirror is a network monitoring tool within Verity that copies traffic from a source port to a destination port. Administrators then connect a PC equipped with Wireshark or other monitoring device to the destination port to analyze the copied traffic.

The **Report → Traffic Mirrors** window provides a Report of all **Traffic Mirrors**.

Creating a Traffic Mirror

MONITORING ACL RESTRICTIONS

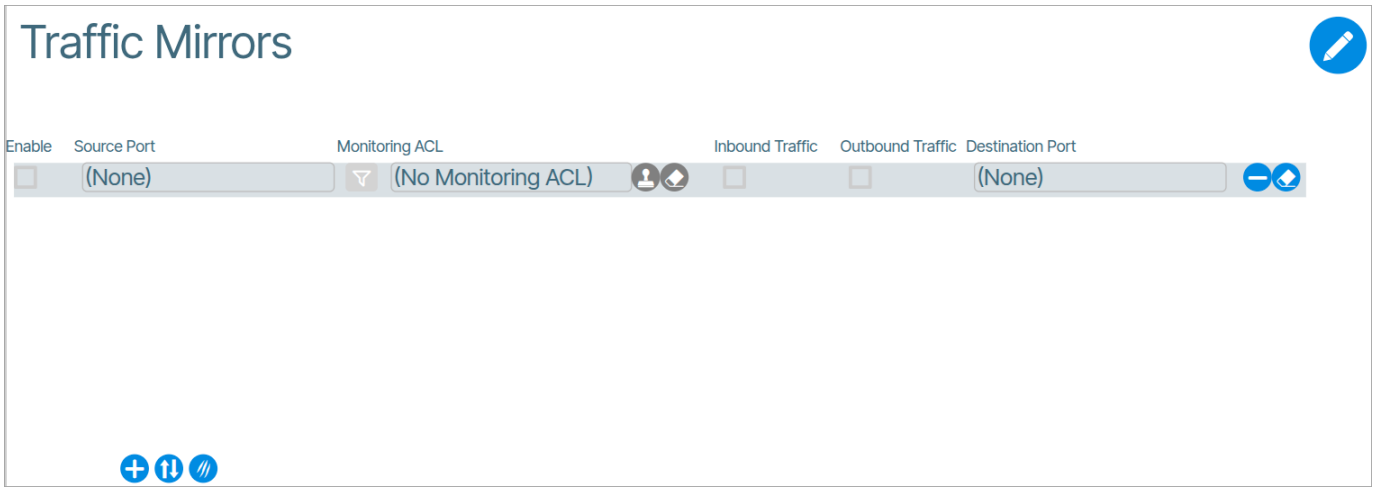
When a Traffic Mirror references a **Monitoring ACL**, the mirror is limited to the Service VLAN traffic selected by that ACL.

The following restrictions apply:

- **All interfaces** is available as a source only when a **Monitoring ACL** is selected.
- Only one enabled Traffic Mirror row on the same switch can use **All interfaces** as the source.
- **Outbound Traffic** is not supported with a **Monitoring ACL**. When a Monitoring ACL is selected, **Outbound Traffic** is disabled and **Inbound Traffic** must remain enabled.

If no Monitoring ACL is selected, choose a specific source port or LAG instead of **All interfaces**.

1. Navigate to **Topology - Network View**
2. Double-click the **Traffic Mirrors** tile on the switch. ( ) This opens the Traffic Mirror window.
3. Click the **Add row to this Switchpoint** button ( ) followed by entering the Source and Destination switch ports, selecting ingress and/or egress traffic, then enable row(s).



Port Traffic Mirroring Indicators

These indicators tell users if a port is used in a traffic mirroring configuration.

Copyright BE Networks 2026



Hovering over the indicator provides more information.



Copyright BE Networks 2026

4.2.10 Packet Broker

A **Packet Broker** is a network device placed between network traffic sources and monitoring tools (or other destinations). It forwards network traffic according to filtering and transformation rules to designated destinations. In Verity, the Packet Broker aggregates network traffic for delivery to multiple destinations. It operates on a dedicated network infrastructure that is physically separated from both the production and management networks. This infrastructure includes packet broker switches, optimized for packet processing rather than general-purpose switching.

Components

Term	Description
Packet Broker Switch (PBS)	A SONiC-based network switch dedicated to aggregating traffic for delivery to multiple analyzers. These switches don't perform typical functions like VLAN switching.
Packet Broker Network (PBN)	A network of interconnected PBS switches, used to forward packets based on Access Control Lists (ACLs).
Packet Broker Ingress Switch (PBIS)	The PBS that connects to the network and receives incoming traffic.
Packet Broker Egress Profile (PBP)	A user-created profile that specifies forwarding rules for traffic, including IP filters for traffic management.
Tap	A hardware device used to passively monitor and copy network traffic without interrupting it.

Functionality

The **Verity Integrated Packet Broker** uses **Packet Broker Switches (PBS)**, which aggregate network traffic for delivery to various **Analyzers**. These PBS switches form a **Packet Broker Network (PBN)**. The PBS switches in this network don’t perform traditional switching functions, such as VLAN switching. Users create **Packet Broker Egress Profiles (PBP)** to define which traffic should be forwarded, applying these profiles to any port in the PBN.

Verity automatically detects the PBS connections and configures the interconnected ports, designating all other ports as input ports. Once **PBPs** are assigned, Verity sets up the packet forwarding rules from ingress to egress.

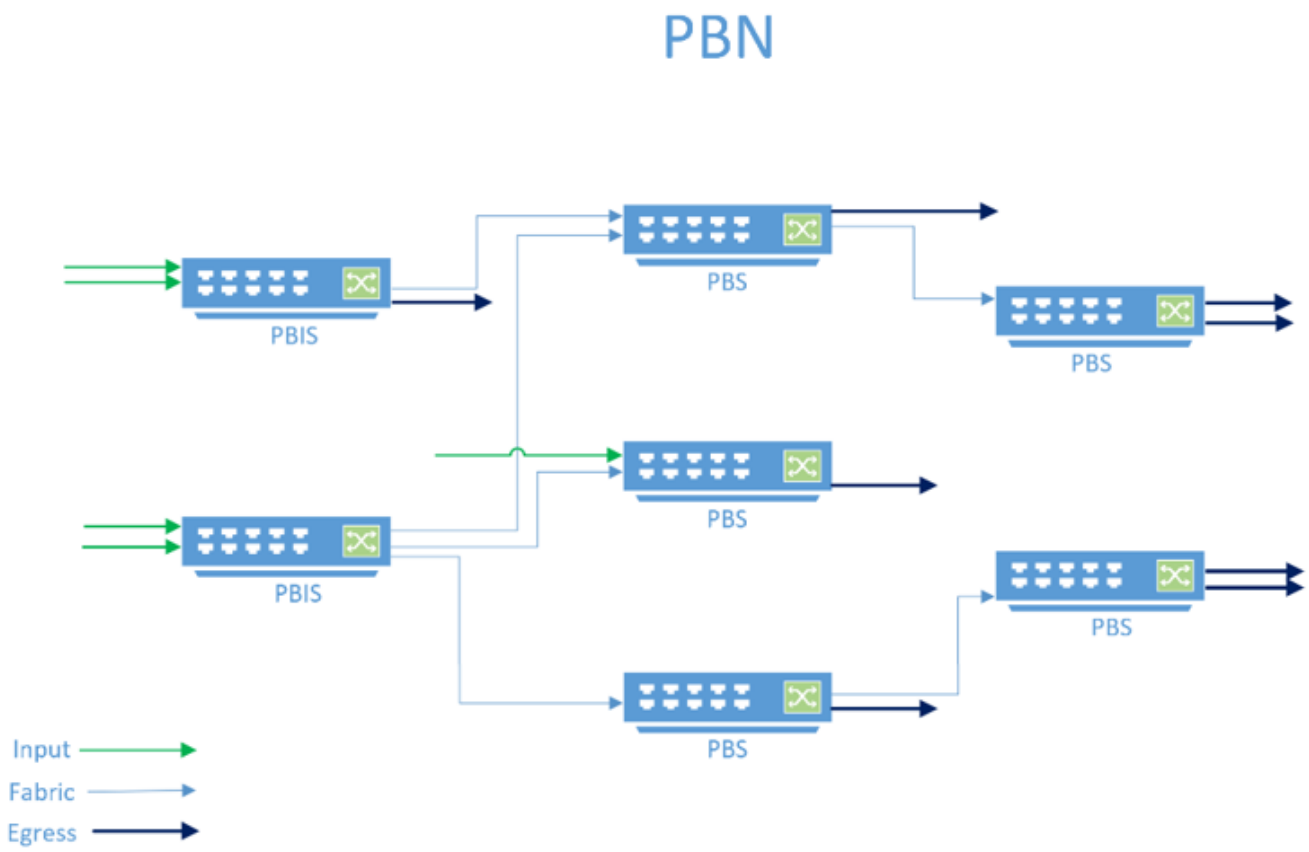
Verity manages Packet Broker Switches utilizing out-of-band ports. The **PBN** is displayed in the **Topology** section of the Verity interface, where PBS switches are shown in left-to-right order, representing the traffic flow. The **PB Ingress Switch (PBIS)** is where traffic begins. The PBN can have multiple PBISs depending on the traffic inputs and load, though PBISs cannot connect to each other.

Each port in the PBN can have a **PBP** assigned, but traffic can only flow from ingress to egress, as shown in the GUI from left to right. When PBS switches are connected, **Link Aggregation Groups (LAGs)** are created automatically for the connected ports. It is the user's responsibility to ensure that the fabric (the ports connecting PBS switches) has the capacity to handle the required traffic.

Architecture

This high-level example diagram shows how traffic flows through the PBN. While there's flexibility of where to place inputs, connections and egress ports, it is recommended to connect input ports to designated PBISs and place egress ports on the lowest tier switches (shown on the right of the diagram and Verity GUI). The network should follow a non-blocking architecture so all input traffic can reach egress switches without congestion.

Copyright BE Networks 2026



PACKET BROKER CALCULATOR

This spreadsheet provides estimates for how many switches are needed based on traffic capturing capacity requirements.

Packet Broker Calculator available here.

Profiles and Filters

The **PBPs** include IP filters that specify which traffic is allowed or denied based on criteria such as protocol, IP address subnet, and port. These filters can be tailored to allow or block specific IP addresses, protocols, and ports, including finer control with **IPv4** or **IPv6** addresses.

UI Overview

In the Verity interface, these sections are used for managing the Packet Broker:

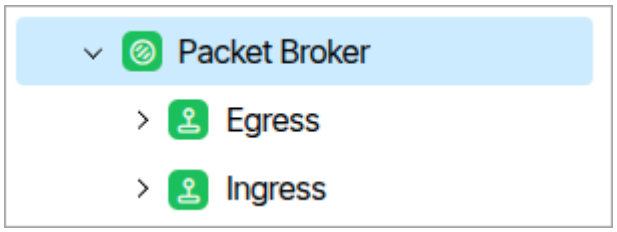
- **Fabric → Packet Broker**
- **Templates → Filters → IPV4 Filters / IPV6 Filters**
- **Templates → Diagnostics → PB Egress Profiles**

FABRIC VIEW

- **Fabric → Packet Broker**

Packet Broker switches are organized as Egress or Ingress.Packet Broker switches are organized as **Egress** or **Ingress**. To add one, select the switch type and click **Add Packet Broker Switch** in the upper right of the window.

Copyright BE Networks 2026

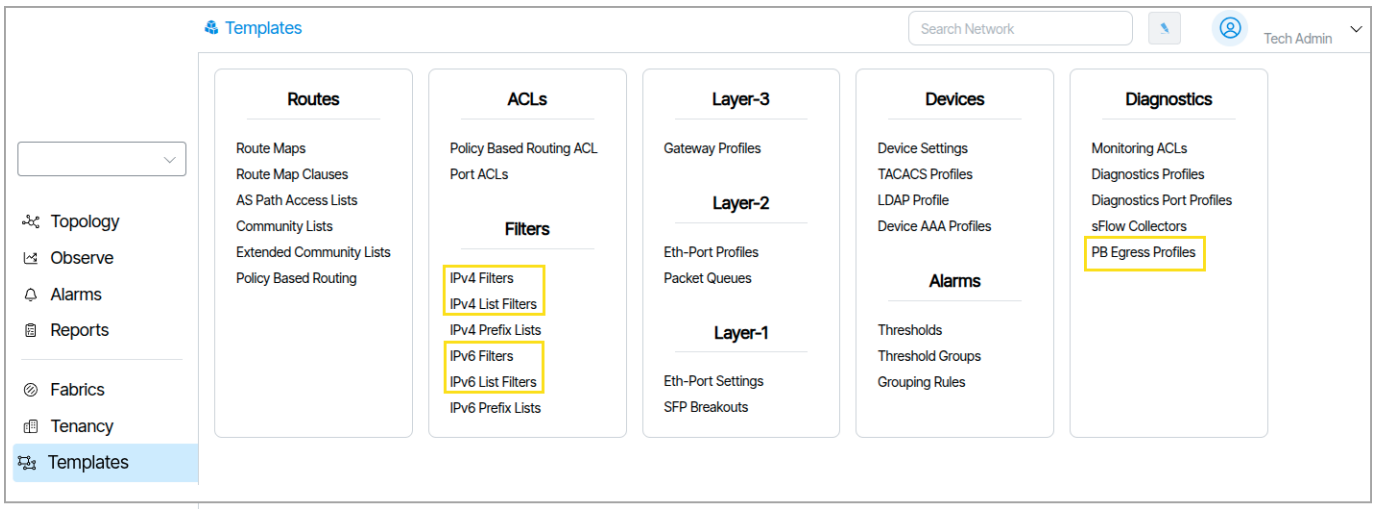


PACKET BROKER TEMPLATES

The Template navigation menu contains these options used to configure Packet Broker.

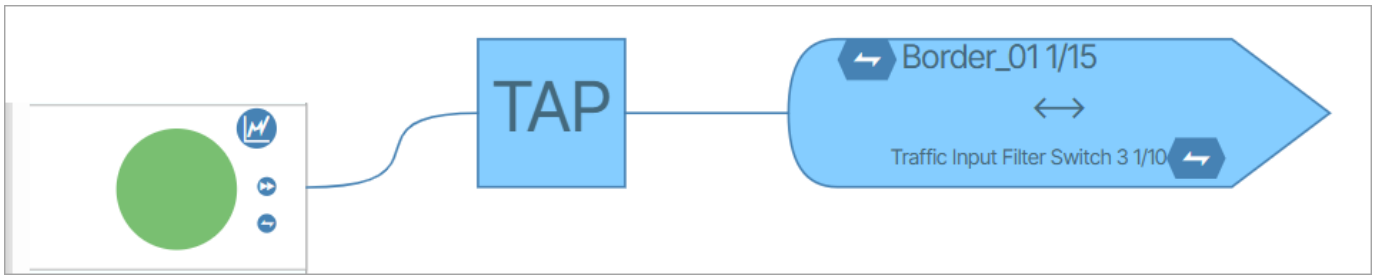
- **Templates → Filters → IPV4 Filters / IPV6 Filters / IPv4 List Filters / IPv6 List Filters**
- **Templates → Diagnostics → PB Egress Profiles**

These tools let users create filter rules and group them into collections called **PB Egress Profiles**, which are then assigned to specific ports in the PBN.



TAPs

A **TAP** (Test Access Point) is a hardware device inserted into the network cabling to passively copy network traffic. The TAP icon in the user interface represents the point where network traffic is accessed and copied for the Packet Broker network.



Creating & Applying Filters

Filtering determines whether selected network data is forwarded to a destination device or blocked.

1. **Determine the Traffic to be Screened**: Identify IP type (IPv4 or IPv6), protocol by name (ip/tcp/udp/icmp only) or IANI number, bidirectional or unidirectional, source and/or destination IP or IP range, source and/or destination ethernet port(s) or range of ethernet ports with appropriate port operators (if applicable). Determine if filter is to Deny and/or Permit traffic types.
2. **Create Filters**: Go to **Templates → Filters** and create filters based on the type of traffic identified in step
3. **Apply the PB Egress Profile to Ports**: Assign the created PB Egress Profiles to device ports.

Copyright BE Networks 2026

APPLY FILTER RULES TO A RANGE OF IP ADDRESSES

1. Select either IPv4 Filters or IPv6 Filters, depending on the address type. 
2. Click the **Add** button, name the filter and click the **Create Filter** button 
3. In the filter settings, edit the rules and apply the changes. 

Filter Rules

Filter	Function
Protocol	IP protocol by name (ip, tcp, udp, or icmp, only) or IP Protocol by IANI number (0-255).
Bidirectional	Bidirectional packets if selected, else unidirectional.
Source IP	IPv4 or IPv6 packet source address. Use IP mask for range of IPs.
Source Port Operator	This field determines which match operation will be applied to TCP/UDP port(s). The choices are equal, greater-than, less-than or range.
Source Port 1	This field is used for equal, greater-than or less-than TCP/UDP port value in match operation. This field is also used for the lower value in the range port match operation.
Source Port 2	This field will only be used in the range TCP/UDP port value operation to define the top value in the range
Destination IP	IPv4 or IPv6 packet source address. Use IP mask for range of IPs.
Destination Port Operator	This field determines which match operation will be applied to TCP/UDP ports. The choices are equal, greater-than, less-than or range.
Destination Port 1	This field is used for equal, greater-than or less-than TCP/UDP port value in match operation. This field is also used for the lower value in the range port match operation.
Destination Port 2	his field will only be used in the range TCP/UDP port value match operation to define the top value in the range.

FILTERING INDIVIDUAL ADDRESSES AS LISTS

To filter individual IP addresses or lists of individual IP addresses, use an **IPv4 List Filter** or **IPv6 List Filter** and enter the addresses in the provided field. A comma must be used between each IP address. There are no spaces used in the lists. 

1. Navigate to **Packet Broker Templates**.
2. Create an **PB Egress Profile** and enable it. 
3. Edit the **PB Egress Profile** by assigning filters to any of the the four available options:

- IPv4 Deny
- IPv6 Deny
- IPv4 Permit
- IPv6 Permit

4. Check the Enabled box to activate the filter(s). 
5. Use the plus button at the bottom-right of the list to create additional rows for assigning multiple filters. 
6. When in Edit Mode save changes by clicking the Save button  .

APPLYING PB EGRESS PROFILE TO DEVICE PORT

1. Once the filters are configured and saved, apply the PB Egress Profile to a device port.
2. Double-click on the device port in the Packet Broker Topology view.

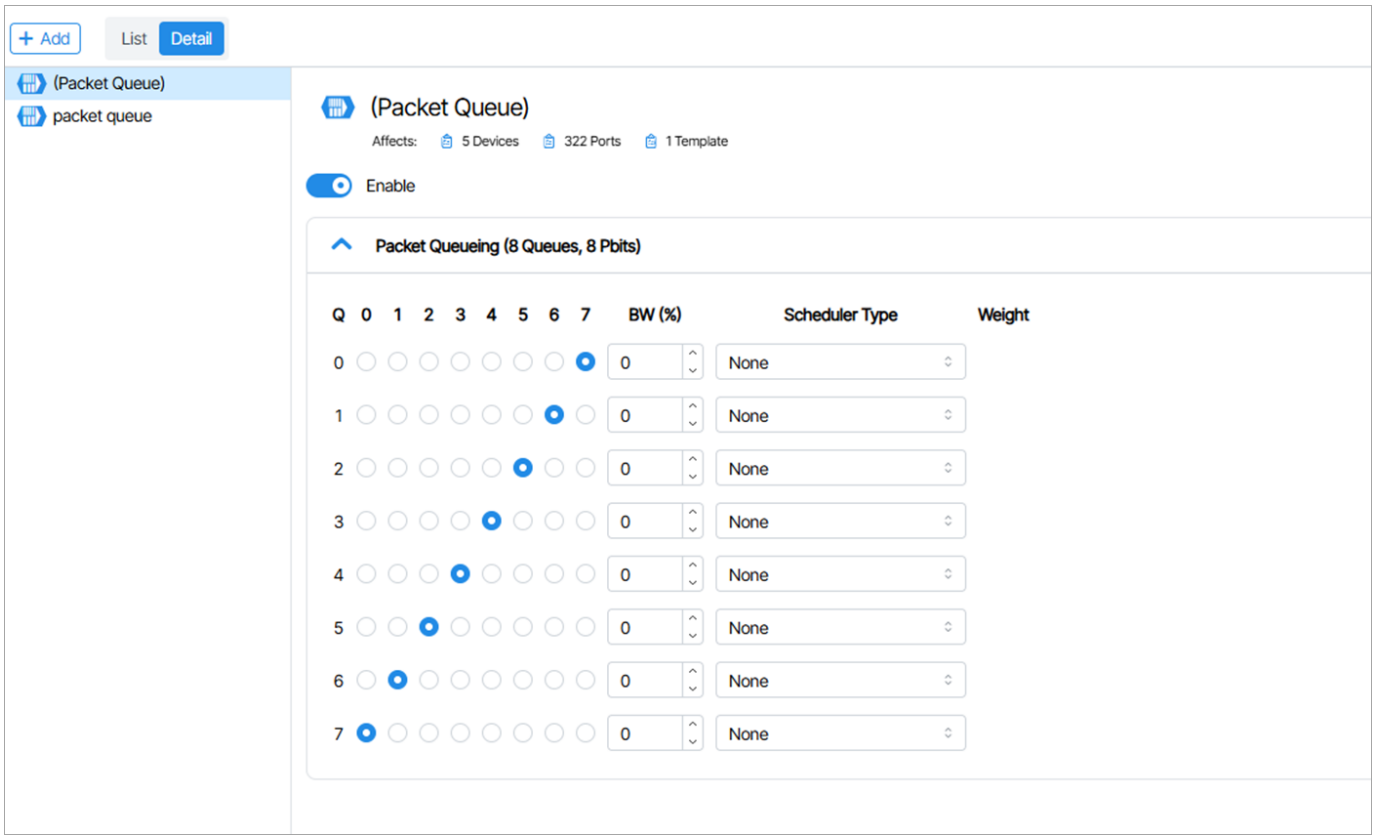
Copyright BE Networks 2026

3. Select the appropriate PB Egress Profile from the dropdown menu. 

Packet Queuing

This feature lets you set the priority for packets depending on bandwidth and is available at **Templates → Layer-2 → Packet Queues**  .

The window that appears lets you make your changes. Here you set the priority for p-bits and set the bandwidth (**BW**) for each queue as a percentage. The percentage is the maximum average percentage of the port output bandwidth that packets in that queue will be allowed to consume. The value cannot exceed 100 and a value of 0 means no limit is set.



Copyright BE Networks 2026

4.3 Tenancy

4.3.1 Multitenancy Overview

Multitenant Fabric Design

Verity provides turnkey data center networking using the popular Leaf-Spine Clos Fabric design. Multiple fabrics can be managed by a single Verity instance. Fabrics can be 3-Stage, 5-Stage or Collapsed Core designs. These networks are all built and operated with Layer-2 ethernet switches and Layer-3 EBGP logical routers.

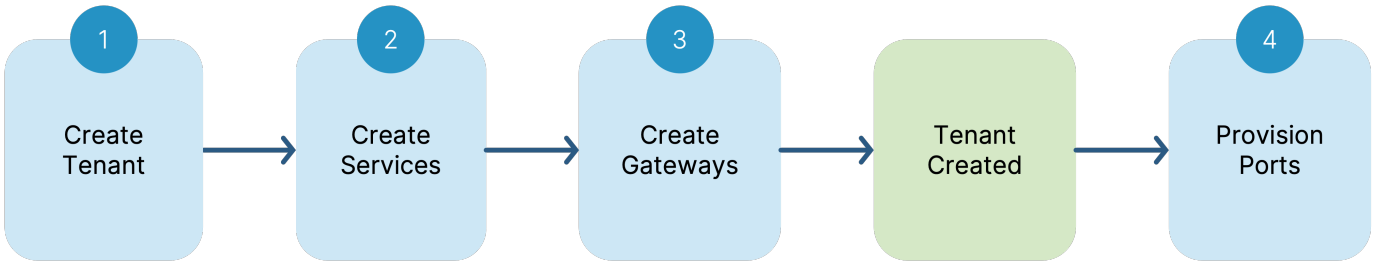
Layer-2 Isolation

Tenants in a Verity system are isolated at Layer-2 with a number of different features and protocols. For bare metal and virtual machine isolation, 802.1q VLANs isolate broadcast domains and create subnetworks on individual switches. This isolation can be extended into a hypervisor host with 802.1q trunks. VLANs can be connected across the Layer-3 routed domain via VXLAN encapsulation and EVPN route advertisements. VLANs are not connected to L3 routers by default, the operator must enable L3 forwarding, DHCP relay, and an anycast gateway which serves as the default gateway for routing purposes.

Layer-3 Isolation

Multiple Layer-2 networks are connected via EBGP virtual routers, and these routers can be isolated from each other using Virtual Routing and Forwarding (VRFs). VRFs are equivalent to a Tenant in a Verity system. Each VRF has its own IP address space, supporting overlapping address ranges. VRFs/Tenants support separate gateways permitting egress from the topology on dedicated or shared L2/L3 interfaces.

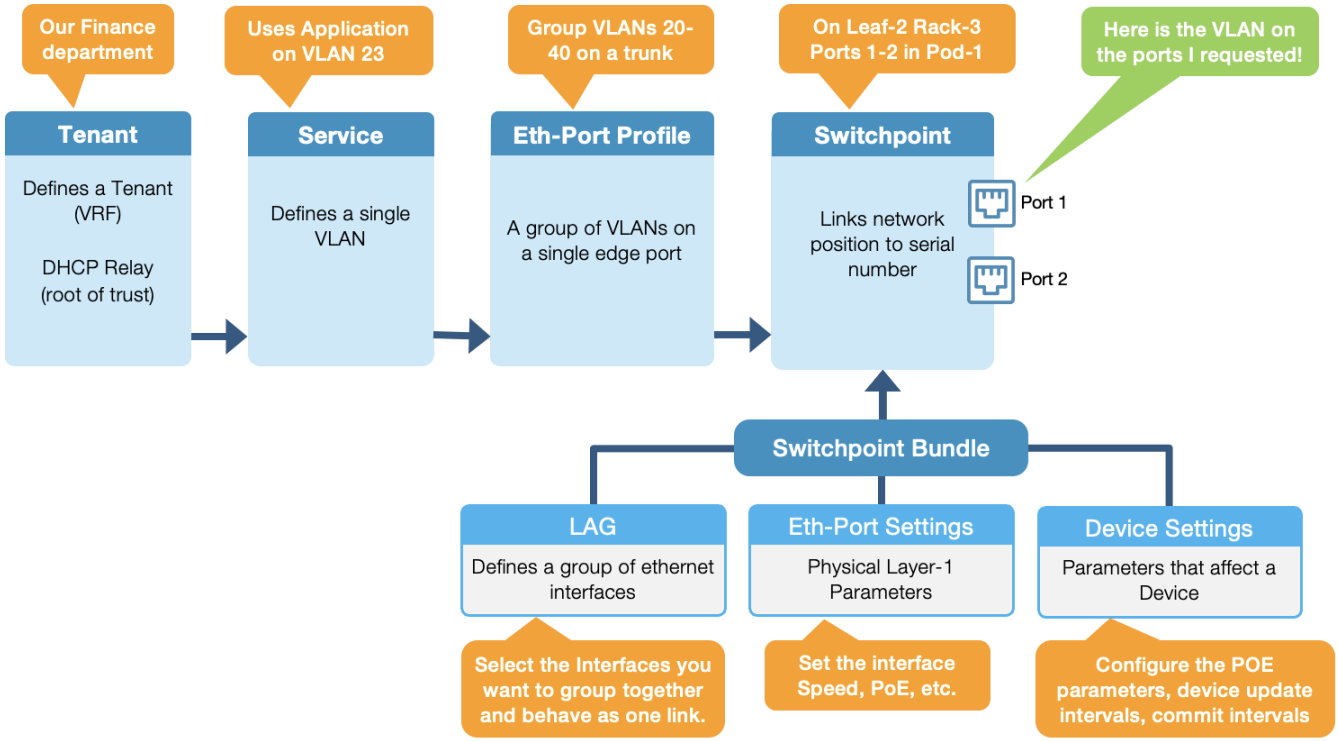
Provisioning Workflow



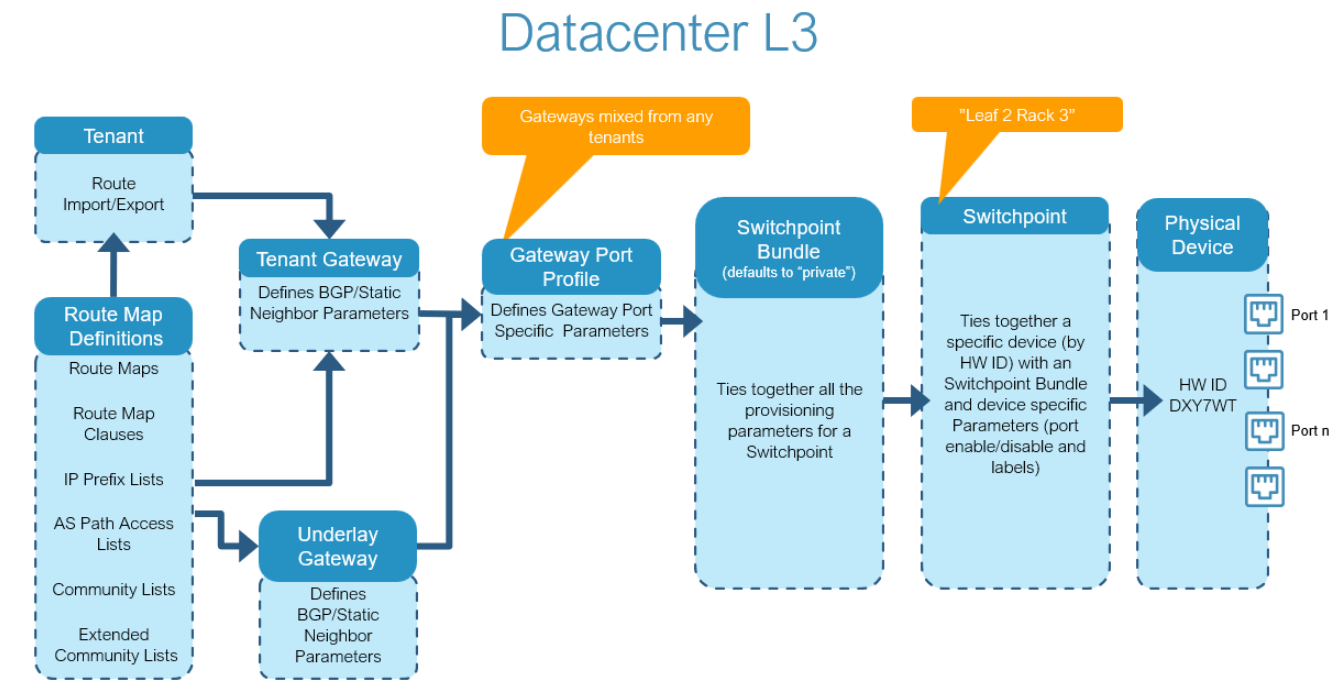
Network provisioning in a Verity system roughly follows this flow. In reality, creating Tenants and Gateways is typically infrequent and once the initial configurations and templates are in place, the operator will most likely be using the Services and Ports objects.

Copyright BE Networks 2026

Layer-2 Provisioning Process (Detailed)



Layer-3 Provisioning Process (Detailed)



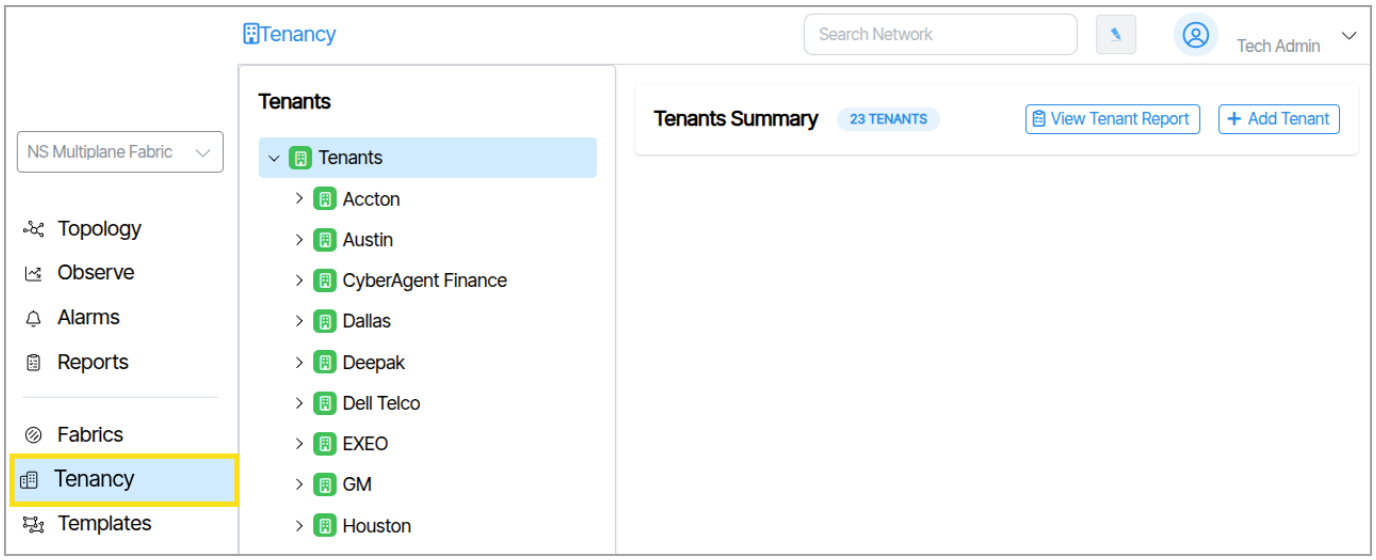
Copyright BE Networks 2026

4.3.2 Tenancy Objects

Network tenancy is controlled through the use of the following objects:

- Tenants
- Services
- Gateways

Tenants, Services, and Gateways are displayed in a collapsible tree view. Each Tenant groups its associated Services and Gateways as child nodes beneath it, and each Service and Gateway parent node contains its own child nodes.  .



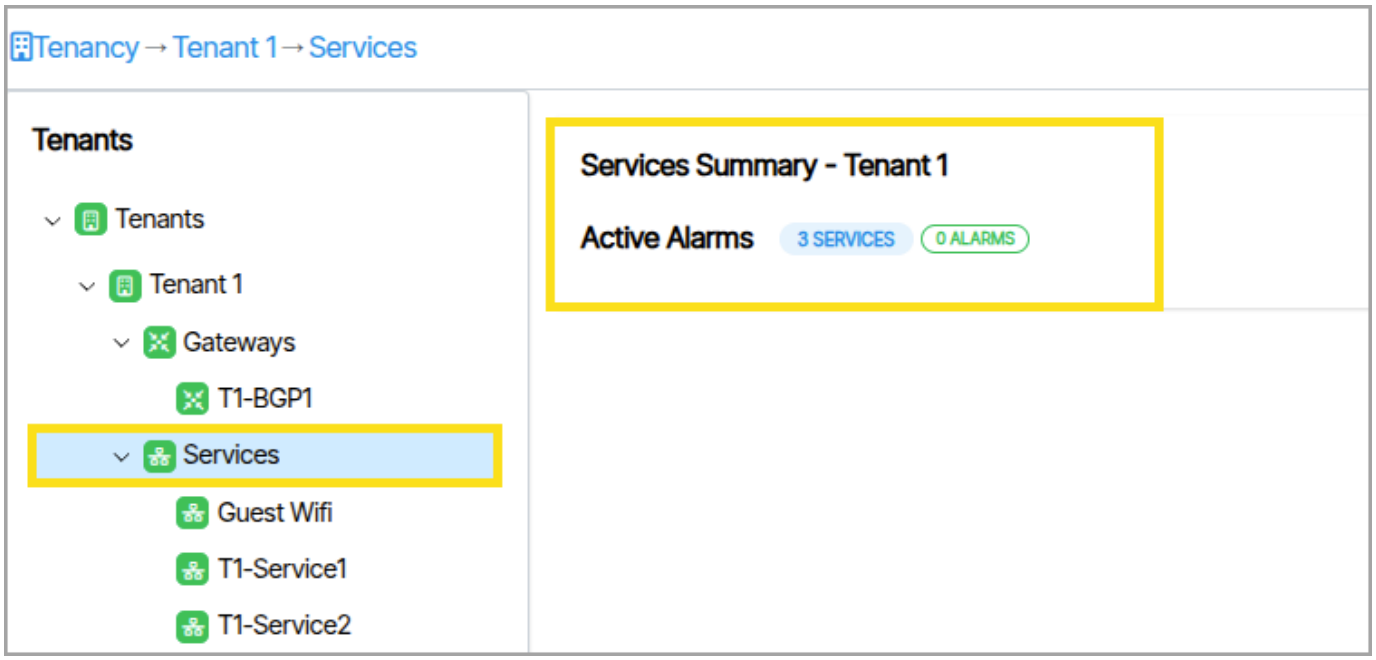
Tenant Tree View Icon Color

Icons display as either green (enabled) or grey (disabled).

Gateway and Service Parent Node Information

Selecting the parent nodes of Gateways or Services displays additional group information that includes the item count and number of active alarms.

Copyright BE Networks 2026



CONFIGURATION

Clicking a Tenant, Service or Gateway reveals the configuration parameters for the selected object.

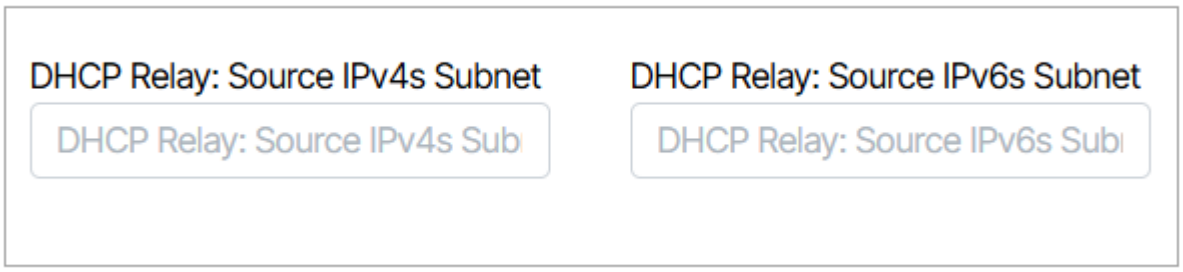
Copyright BE Networks 2026

4.3.3 Tenants

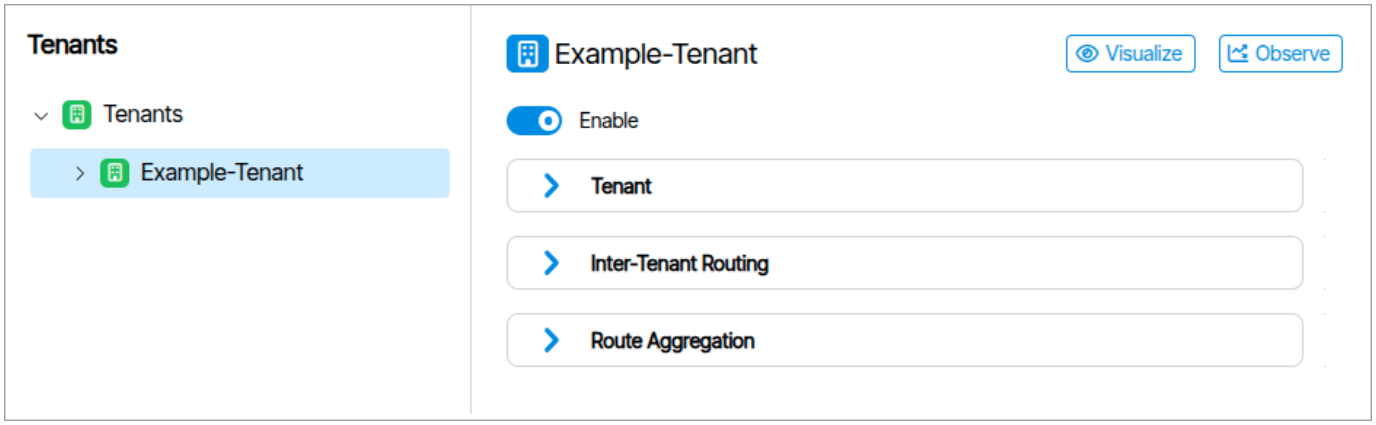
Tenants are the fundamental building block of data center overlay traffic. A Tenant is composed of several elements, including Services, a Layer 3 routing domain, and a BGP gateway to enable communication outside the data center. Proper configuration of Tenants is crucial to setting up a data center deployment. This document offers general instructions on how to create and configure Tenants.

Creating Tenants

1. Navigate to **Tenancy → Tenants**
2. Click the **Add Tenant** button (). Type the name of the Tenant and save it by clicking the **Create Tenant** button.
3. Explore the parameters. 
4. The primary **Layer-3** parameters are auto-generated. This includes **VRF Name**, **Layer-3 VNI**, **Layer-3 VLAN**.
5. Optionally configure the **DHCP Relay Source IPs Subnet** which is used to allocate IP addresses to connected devices within the Tenant (additionally controlled within the **Services** object).



Selecting Tenants



Selecting a Tenant displays the following drop down items. Each contain editable parameters.

- Tenant
- Inter-Tenant Routing
- Route Aggregation

Visualization

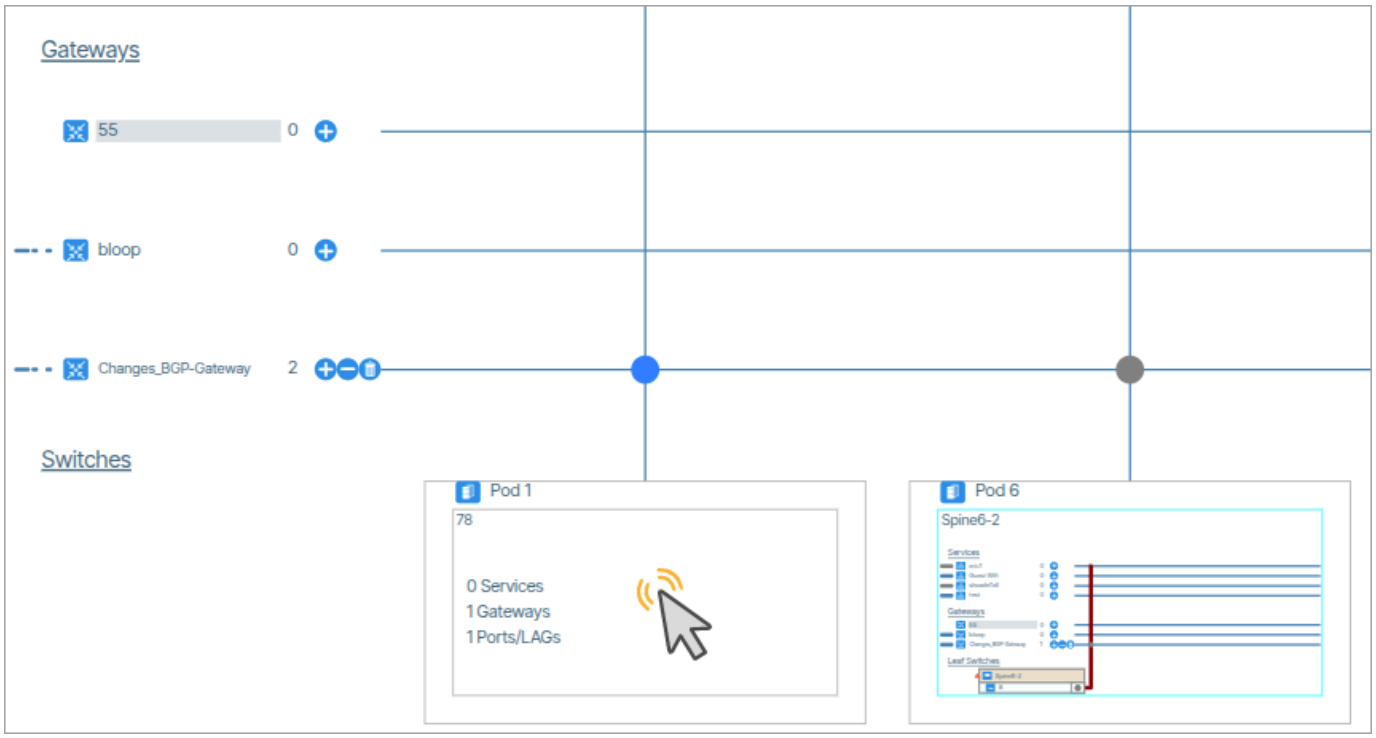
The Visualize button opens a visual interface called **Tenant Slice**.

Copyright BE Networks 2026



The **Tenant Slice** window offers a visual interface to assign services and Tenant BGP Gateways to Leaf ports. However, when assigning a Service to a Leaf port, the Service is not directly assigned to the port itself. Instead, the Service is assigned to an **Eth-port-profile** that is *automatically generated* and assigned to the Leaf port. Eth Port Profiles are collections of one or more Services that you assign to Leaf ports or external LAGs. To summarize, Services are assigned to Eth Port Profiles and Eth Port Profiles are assigned to Leaf Ports/LAGs.

When a Tenant has Services/BGP Gateways provisioned across switches from more than one POD, the data is condensed into visual boxes that encapsulate connections. To view encapsulated connections, double-click the boxes.



TENANT SLICE OPERATIONS

Assign Services to the LAG via Tenant Slice

The conventional way to assign Services to a LAG is via the Tenant Slice window.

1. Select a **Tenant** and click **Visualize**. 
2. Zoom in and select the **Service** you want to assign and click the Add Service button  .
3. In the drop-down menus, find the desired LAG and select it. Click the checkbox button to create the assignment. 
4. When complete, a visual representation expressing the assignment appears. 

Copyright BE Networks 2026

Assign a Service to a Leaf Port

The following is an example of assigning a **Service** to a **Leaf** port.

1. Click the **Add Service** button ( ).
2. While in **Tenant Slice**, select the **Service** you want to use and click its plus button .
3. In the window that appears, use the drop-down menus to select the (switch and) port to assign the **Service** to.
4. When the assignment is complete, the **Tenant Slice** window updates to reflect the new state.

Removing an Assignment

Remove an assignment by clicking either the **Remove Service** ( ) icon or the node connection ( ) icon of the assignment you want to delete.

Assigning BGP Gateways

The mechanics of creating and managing BGP Gateway assignments in Tenant Slice are the same as the previously discussed Service creation and navigation. When creating a new BGP Gateway assignment, the user is prompted to provide an IP address and a subnet mask in the format of this *example*: 192.168.1.0/24

PORTS

The number of Leaf ports used by a collection of Tenant Services and Tenant BGP Gateways.

To manipulate port assignments, you use the add and remove buttons:



To assign port configurations to the Tenant Slice, the initial step involves identifying the **Service** or **BGP Gateway** to be routed. Then, proceed to click on the designated add or remove button ( ) and perform the necessary assignment or removal to the switch port displayed in the **Tenant Slice** window.



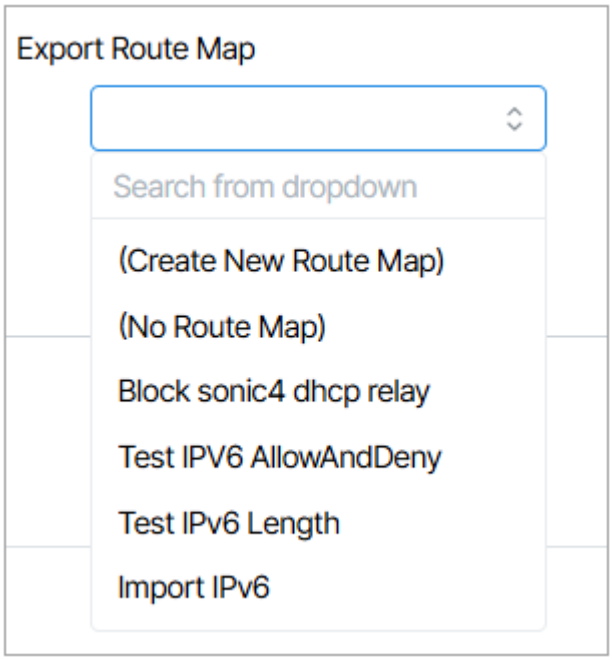
ROUTE MAPS

To create Route Maps, refer to the main Route-Maps section.

Apply Route-Maps to Tenant-to-Tenant Routing

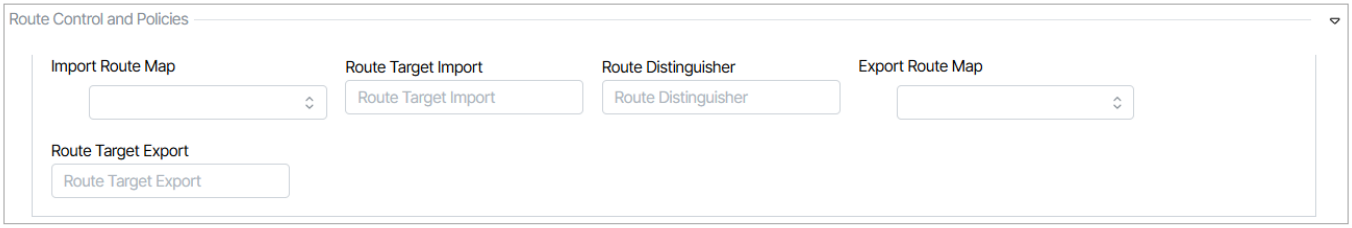
To apply Route-Maps to your Tenant-to-Tenant routing click on the Route Map field and make your selection from the drop down menu.

Copyright BE Networks 2026



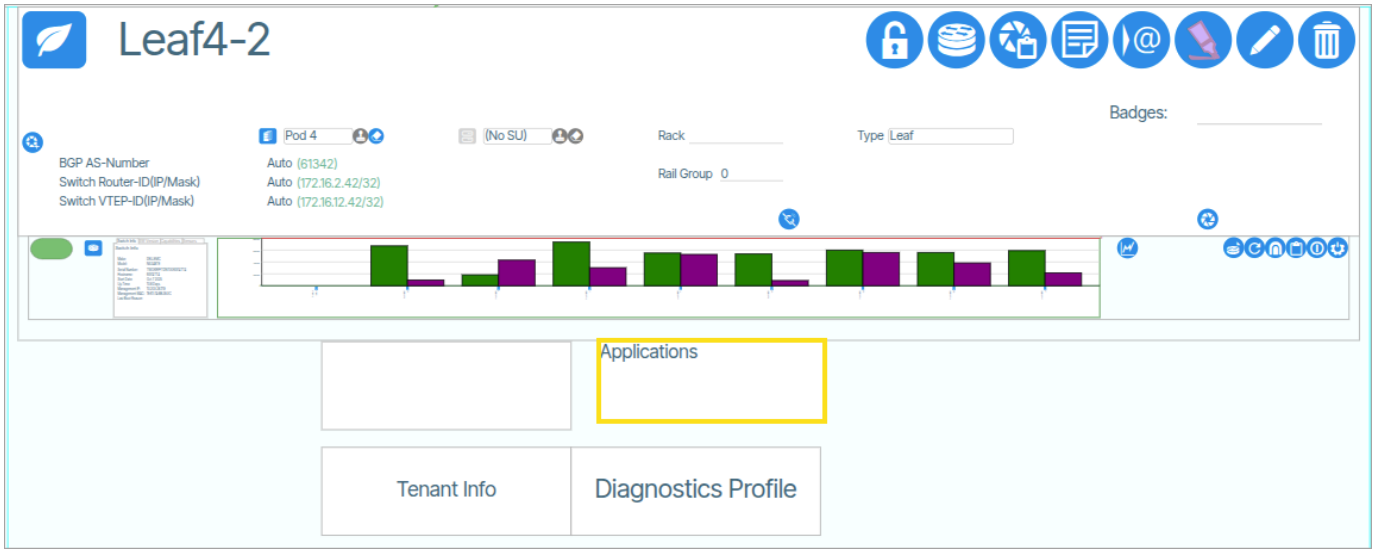
Route Target, Route Distinguisher

To set Route Target or Route Distinguisher values, select the respective field and type in the value.



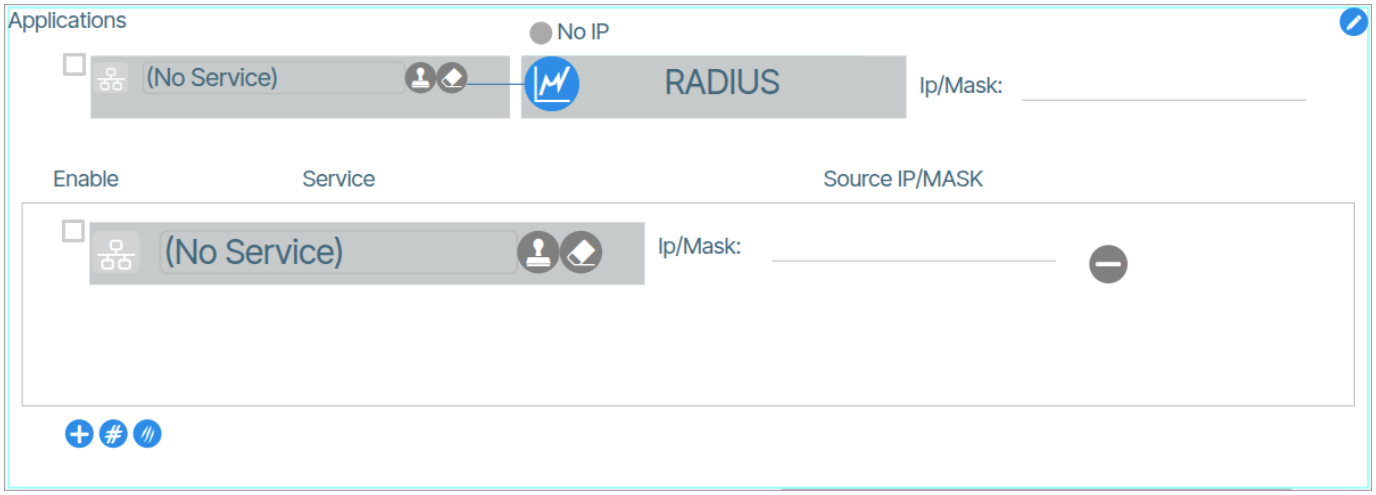
Applications

Applications enable users to select VLAN interfaces and configure IP address assignments for those interfaces on the switch. This feature enables network segmentation and service configuration by allowing users to associate IP addresses with defined Services.



Copyright BE Networks 2026

Each Service represents a logical grouping of VLAN configuration parameters, providing a streamlined approach to network resource allocation and traffic management.



CREATE IP-TO-VLAN MAPPING

1. To perform the IP-to-VLAN mapping you need to edit-enable ( ) in the Applications tile.
2. Select the desired Service from the **Service** drop-down menu.
3. Under the field titled **Source IP/MASK**, type the desired IP/MASK.
4. Enable the change Enable and save the edit ( ).

Additional Resources

- Multitenancy

Copyright BE Networks 2026

Anycast Service Tenant VLAN VXLAN

4.3.4 Services

Services are Layer-2 broadcast domains (VLANs) that are optionally routed through the Layer-3 fabric as VXLAN frames. Services enable devices connected to different parts of the network to appear to be in the same local segment. Layer-3 addressing and routing can be optionally enabled for each service to support IP packet forwarding as well as DHCP relay services. A service is linked to a tenant for further isolation of network reachability.

Creating a Service

1. From the Main Navigation select Tenancy. Under the Tenants tree select the desired Tenant and unfold it to reveal the Services option beneath it. 
2. Select Services. Click the **Add** a new Service button  . Type a name in the field that appears and click **Create Service**
3. The Service Enable switch defaults to *off*. The Tenant field shows which tenant the Service is assigned to 

Setting	Value	Description
Layer-2 VLAN	Required	The 802.1q VLAN tag to be used on the provisioned switch ports.
Layer-2 VNI	(Autoassigned)	Specifies the VXLAN VNID number if the user wishes to override the automatic assignment.
Anycast Address	Optional	Comma-separated lists of anycast gateway addresses for IPv4 and IPv6 subnets. Each anycast address serves as the default gateway for a different subnet within the same VLAN, enabling Layer-3 routing and reachability across multiple subnets.
DHCP Relay	Optional	Enables DHCP packet forwarding to a central DHCP server(s) for address requests.
MTU	Optional	Maximum Transmission Unit for the Layer-2 segment. This is either *automatically generated* or *not* required.

Note

You must set the DHCP source address range in the Tenant Layer 3 section described below in order to support DHCP Relay. The DHCP Source address range is used by Verity to create loopback addresses within each Leaf switch to serve as the source address for the DHCP requests for the tenant when it is configured on the switch)

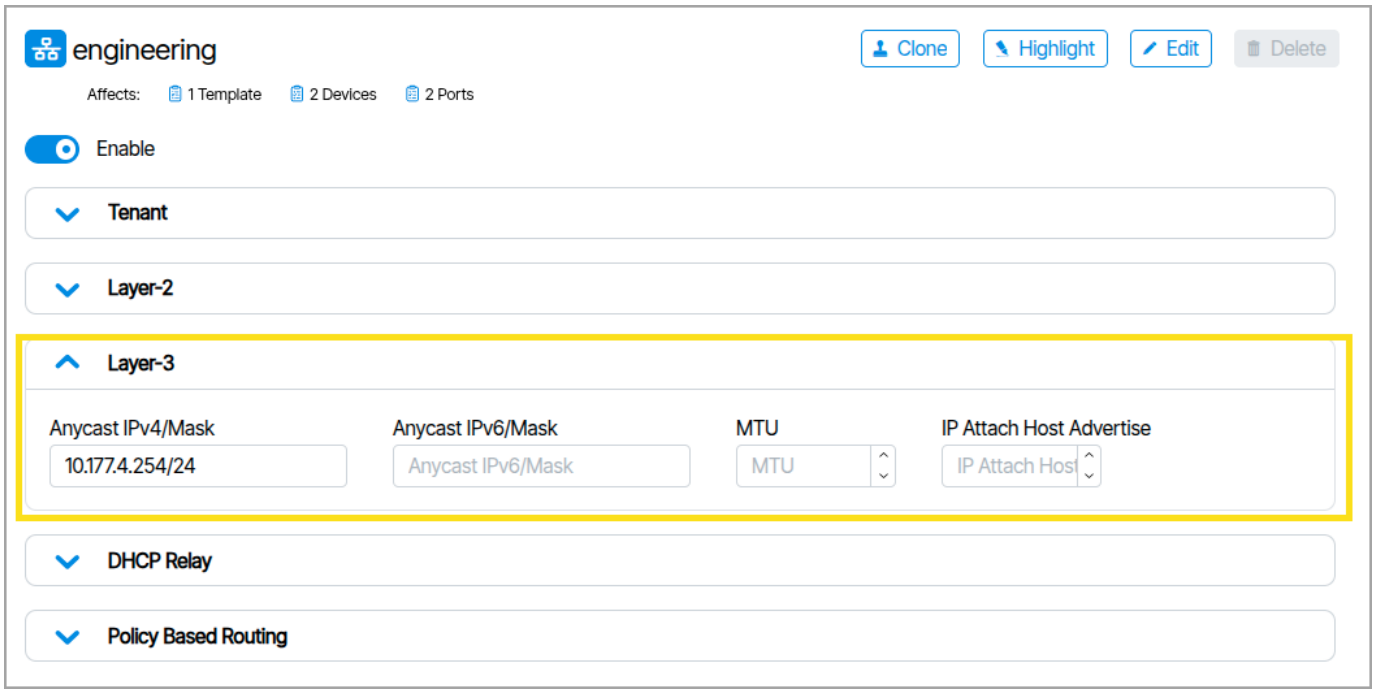
1. Click **Enable** to enable the Service. This will place the Service into **Edit Mode**. 
2. Fill out any required fields.
3. Click Save. 

Enable Layer-3 Anycast IP/Mask on a Service

Note

Prior to enabling Tenant-to-Tenant routing your chosen Tenant must have a Service with a Layer-3 Anycast IP/Mask field set.

Copyright BE Networks 2026



1. Select a Service.
2. In the **Anycast IP/Mask** form field type the static anycast gateway address for this service.
3. Click Save. 

Anycast IPv4/Mask A comma-separated list of static IPv4 anycast gateway addresses assigned to this service.

Anycast IPv6/Mask A comma-separated list of static IPv6 anycast gateway addresses assigned to this service.

MTU The Maximum Transmission Unit defines the largest packet size the switch will forward without fragmentation on this service. Mismatched MTU values within the same VLAN can cause silent packet drops, particularly for larger frames common in storage or virtualization traffic.

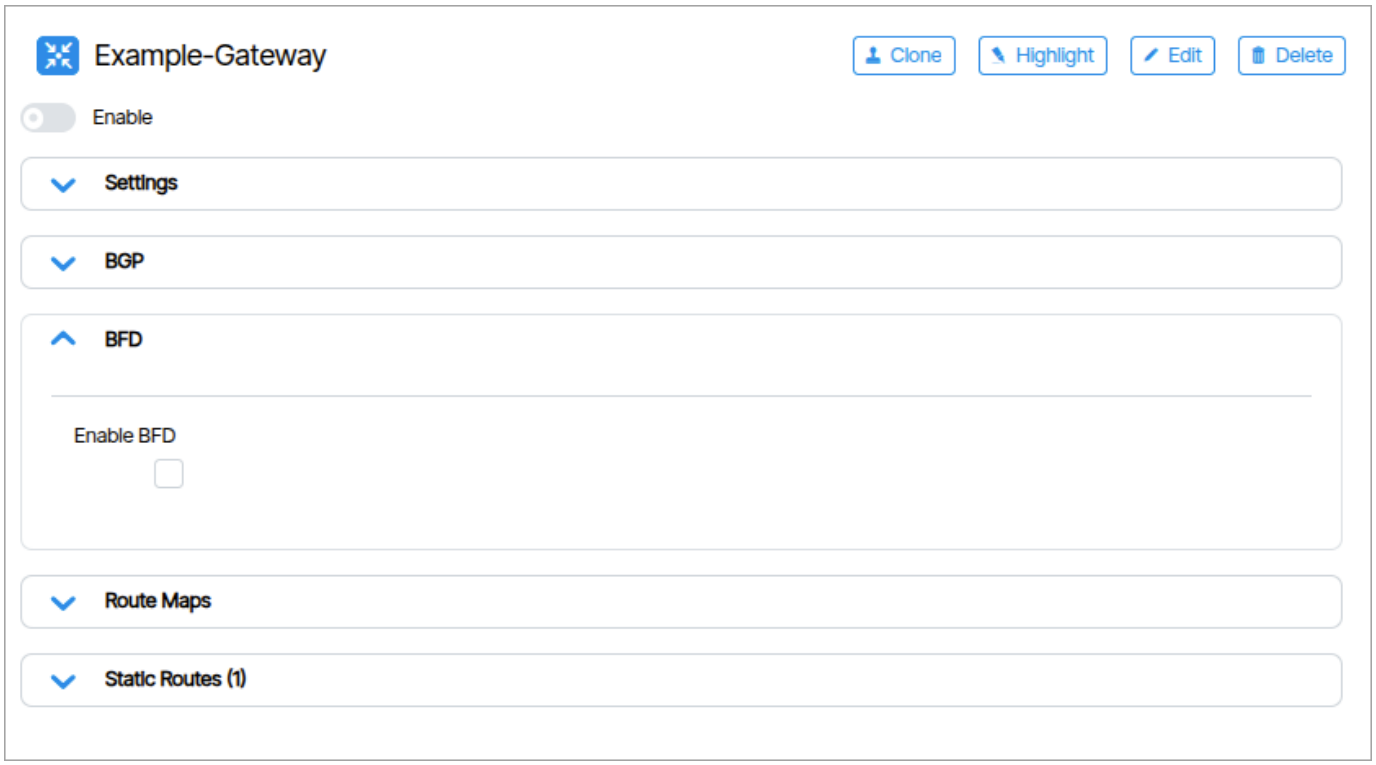
IP Attach Host Advertise Redistributes dynamically learned /32 host routes from attached ports into BGP, making locally discovered hosts reachable across the fabric. Enter a value between 0–250 to set the admin distance for these redistributed routes, where lower values give them higher priority; leave blank to disable the feature.

Copyright BE Networks 2026

BFD Gateway Gateway Modes Gateways

4.3.5 Gateways

Gateways enable traffic flow in and out of a network or tenant to a device external to the Verity managed network.



Gateway Modes

Gateways can be one of several different modes. The **Gateway Mode** ( ) field defines how the software handles and routes traffic within the network. It determines whether routing decisions are based on a statically defined default gateway, static routes, traditional BGP routing, or new enhancements to BGP to support dynamic neighbors .

- **Default Route**: A single static default route.
- **Static Route**: Uses manually configured static routes, no dynamic routing.
- **Static BGP**: Establishes BGP sessions with manually configured BGP neighbors. The BGP neighbors are explicitly defined by their IP addresses, and no dynamic peer discovery occurs within the subnet. The user must specify BGP Neighbor IP address and AS numbers at a minimum. Multi-hop peering is supported.
- **Dynamic BGP**: BGP neighbors are dynamically discovered within a subnet, leading to the creation of multiple BGP sessions from a single tenant gateway object.
- **IPV6 Unnumbered**: Establish BGP peering over a link without assigning unique IPv6 addresses to the physical interface, using the interface's automatically generated link-local IPv6 address (FE80::/10) as the BGP neighbor address.

Creating a Gateway

1. From the Main Navigation select Tenancy. Under the Tenants collapsible tree select the desired Tenant and unfold it to reveal the Gateways option beneath it.
2. Select Gateways.
3. Click the **Add Gateway** button.

Copyright BE Networks 2026

4. Type a name in the field that appears and click **Add Gateway**.
5. If the Tenant is going to egress from the network on an 802.1q tagged interface, the VLAN should be specified in the **Egress VLAN** field ().
6. The Gateway Enable switch defaults to off. Click it to enable the gateway.
7. Click the **Save** button.

USING ANYCAST GATEWAY INTERFACE FOR PEERING

When **Anycast Gateway Mode** is enabled in the Gateway Profile, the **Source IP Address** field in the BGP configuration must be populated. This address is used as the local BGP peering address presented to neighboring routers.

The Anycast Gateway address serves as the default gateway for hosts within the associated subnet. The Source IP Address must be assigned from a different subnet than the Anycast Gateway subnet.

The system automatically creates a loopback interface and assigns the Source IP Address to that loopback. Neighboring routers should use this Source IP Address as the BGP neighbor address when establishing the BGP session.

There should be no IP address assigned to this type of gateway when referenced in a Gateway Profile.

BFD CONFIGURATION (OPTIONAL)

BFD (Bidirectional Forwarding Detection) is a network protocol used to detect faults in the forwarding path between two routers or network devices. It uses heartbeat-timed messages that are exchanged between devices at regular intervals to monitor the status of the link. 

- **BFD Receive Interval**: Configure the minimum interval during which the system can receive BFD control packets.
- **BFD Transmission Interval**: Configure the minimum interval during which the system can send BFD control packets.
- **BFD Detect Multiplier**: Configure detection multiplier to determine packet loss.
- **BFD Multihop**: Enable BFD Multi-Hop for Neighbor. This feature is used to detect failures in the forwarding path between BGP peers.

  In a multi-hop BGP session, there are intermediate switches between the two devices where the session is configured. Instead of the BFD packets traveling directly from switch "A" (where the BGP session originates) to switch "B" (where it terminates), they pass through several intermediate switches along the route.

  Enabling multi-hop BFD is optional, as it depends on the specific requirements of the BGP session. The operator must assess whether sub-second BGP session failure detection is necessary. If such detection is not critical, or if enabling it might result in false positives, it may be better to leave this feature disabled.

Related Information

Gateway Profiles

Copyright BE Networks 2026

4.3.6 Inter-Tenant Routing

Single Tenant-to-Tenant Routing

Before you perform this feature, the selected Tenants must enable their Layer-3 Anycast IP/Mask on a service.

1. Click **Tenancy**.
2. Click your chosen **Tenant** from the tree.
3. View the section titled **Inter-Tenant Routing**.

You create Tenant-to-Tenant associations, set Route Maps, and apply Route Target and Distinguisher settings here. To create an inter-tenant connection, you first need to select a Tenant to connect with. In the Imports field, select a Tenant of your choice and click the **Enable** checkbox next to it. To add more Tenants, repeat the Import process.

BIDIRECTIONAL TENANT-TO-TENANT ROUTING

To configure two-way Tenant Routing, complete the above steps to enable single Tenant-to-Tenant Routing.

1. Under the **Inter-Tenant Routing** section, click the Object button next to the selected Tenant.
2. In the Tenant window that appears, click **Inter-Tenant Routing** and assign the Tenant from the previous window. Click **Enable** to confirm the selection.

When you view the Tenant-Tenant Routing for bidirectional Tenants, both Tenants display their connection to one another in their Imports window.

TENANT-TO-TENANT ROUTING WITH ROUTE MAPS

You need to create Route Maps before you apply them to Tenant-to-Tenant routing. Creating Route Maps is a separate process. Additionally, applying Route Maps to Tenant-to-Tenant routing requires that you already configured Tenant-to-Tenant routing.

1. Click **Tenancy** and click the Tenant you want to work with.
2. View the **Inter-Tenant Routing** section.
3. Set the values of the following five form fields as needed for your configuration:

- Route Distinguisher
- Route Target Import
- Import Route Map
- Route Target Export
- Export Route Map

Copyright BE Networks 2026

4.3.7 Inter-Site Routing

Inter-POD Gateways

This feature enables underlay BGP connections between PODS in lieu of using the Super Spine layer. The typical use case is for multiple PODs ( ) managed by a single instance of Verity system that are separate geographical sites.It is recommended that the Leaf switches used for the BGP connections between PODs are not members of an MCLAG pair.

CREATING SITE GATEWAYS

The BGP connections are created using **Fabric Gateway** objects found in **Fabric → Gateways → Gateway**.

A **Fabric Gateway** is required for each end of the BGP sessions planned between the sites. For high availability this requires four site gateways to be created. Each Gateway object requires the Neighbor AS number of the Leaf on the other side of the planned connection, and the IP address that is used on the Gateway Port Profile on the neighbor's switch. Additionally, the option called "Fabric Interconnect" should be set.

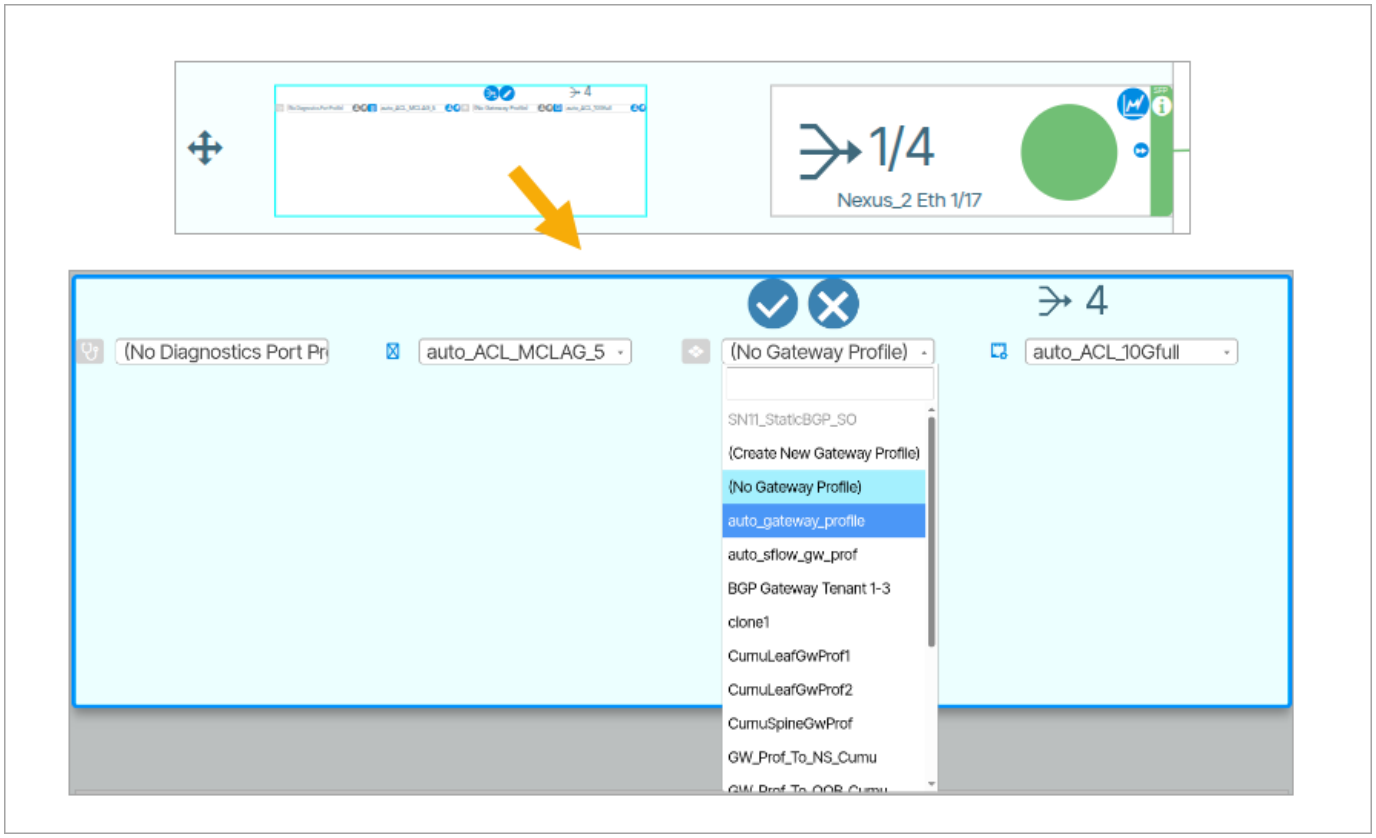
CREATING GATEWAY PROFILES

Go to **Templates/Gateway Profiles** to create the needed Gateway Profiles that are used to apply to either ports directly, or indirectly through LAG objects.

Four profiles are required. Each profile references one of the Gateway objects and contains the IP address assigned to the interface that terminates the BGP session on the switch where it will be applied. This IP address communicates with the neighbor IP address specified in the Gateway object.

ADDING A GATEWAY PROFILE TO A LAG

The Gateway profile can be assigned to a single port or to a LAG. In this example, LAGs are used. Set the Gateway Profile in the **Gateway Profile for Eth** field on the device port.



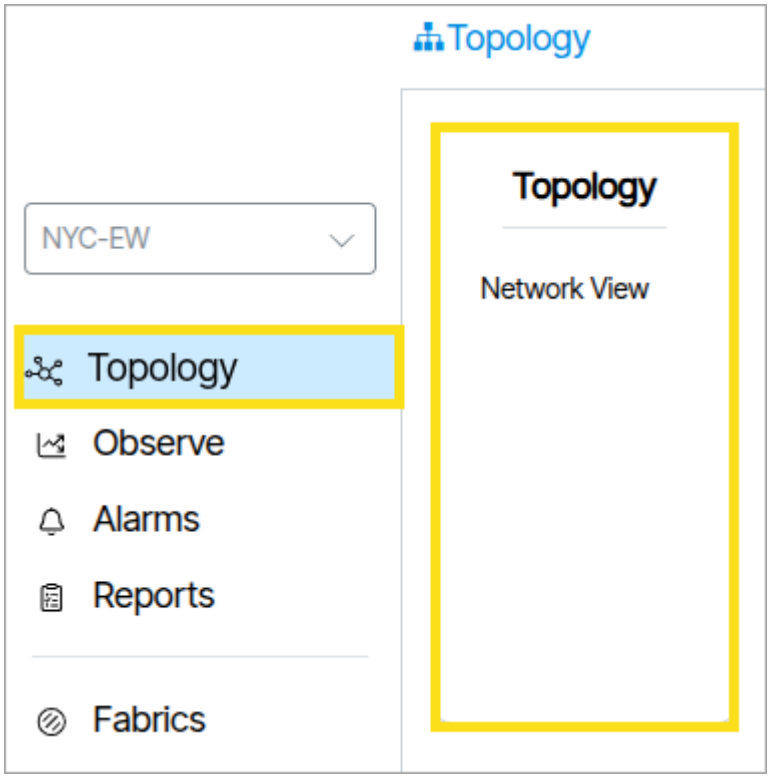
Copyright BE Networks 2026

4.4 Topology

4.4.1 Topology Overview

The Topology section contains navigation links to the following sections and tools:

- Network View
- Management Network

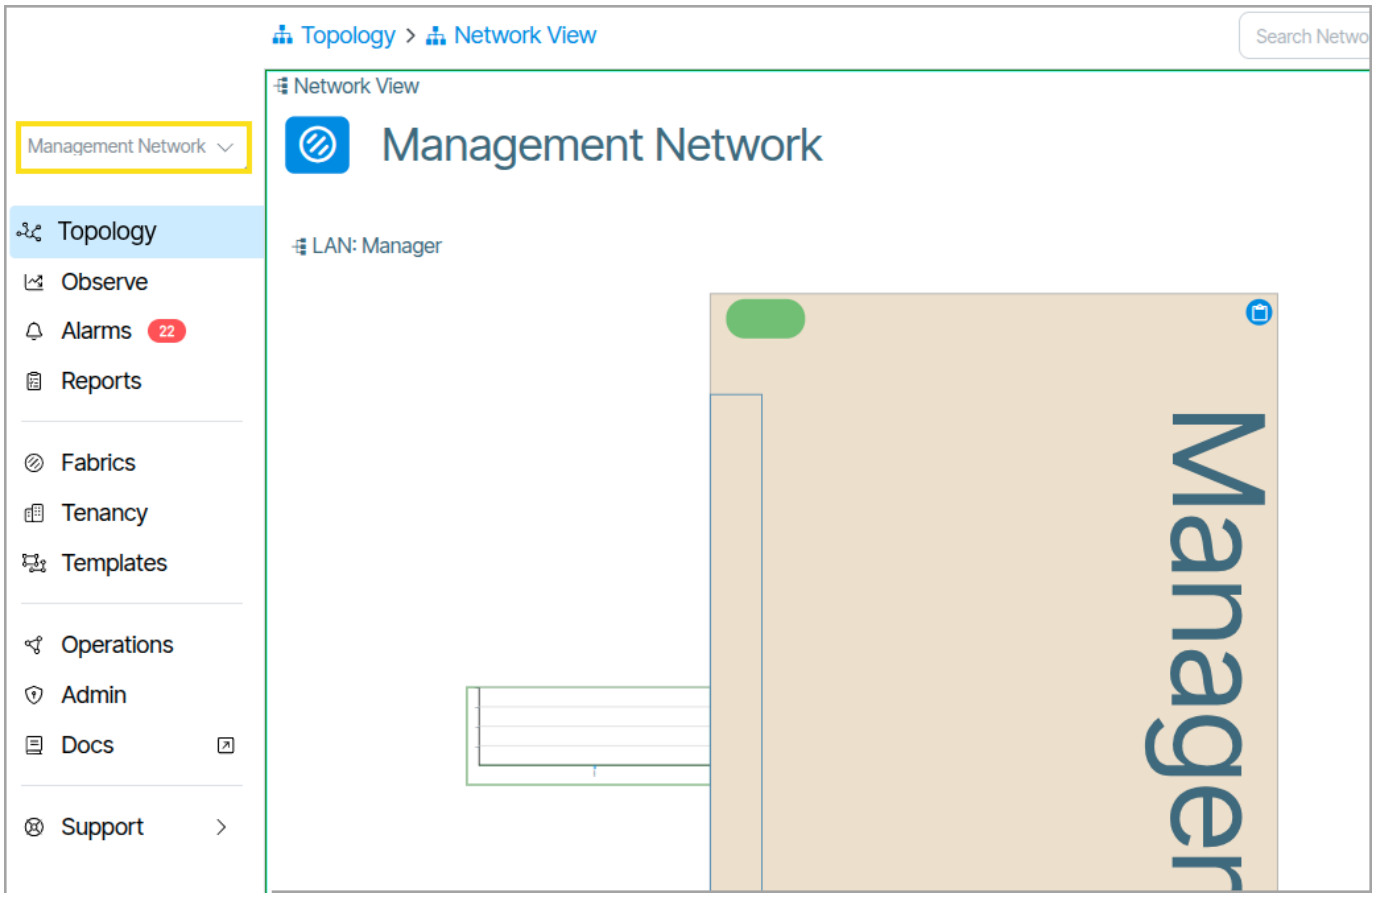


This section is used to visualize the network (Network View).

Management Network

Management Network opens from the fabric search and displays a Management plane visualization.

Copyright BE Networks 2026

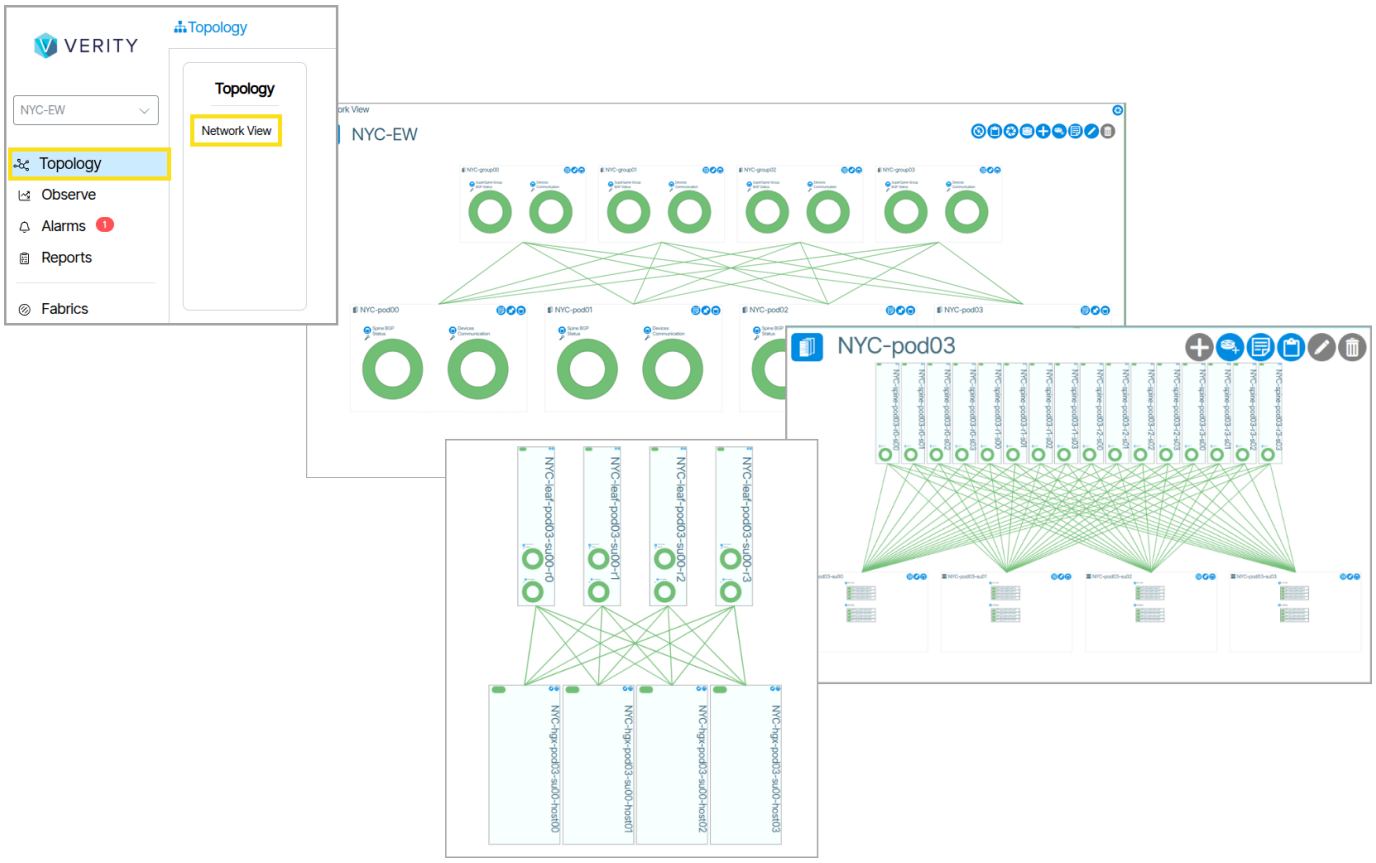


Copyright BE Networks 2026

4.4.2 Network View

The Network View window resides inside **Topology** and displays virtual devices and connections that represent hardware devices in the physical network configuration. The **Network View** is designed to behave like Google Maps in that it uses a zoomable user interface. Panning and zooming allows for the viewing of sections in more detail. The more you zoom in, the more context is shown. Zooming in and out is performed by mouse wheel scrolling or double clicking to focus on a section.

Pan - Hold-click on any point of the map and move the mouse. Keyboard arrow keys also support panning. **Zoom** - Double-click on an object to go to a defined view. You can also drag and scroll the page to navigate.

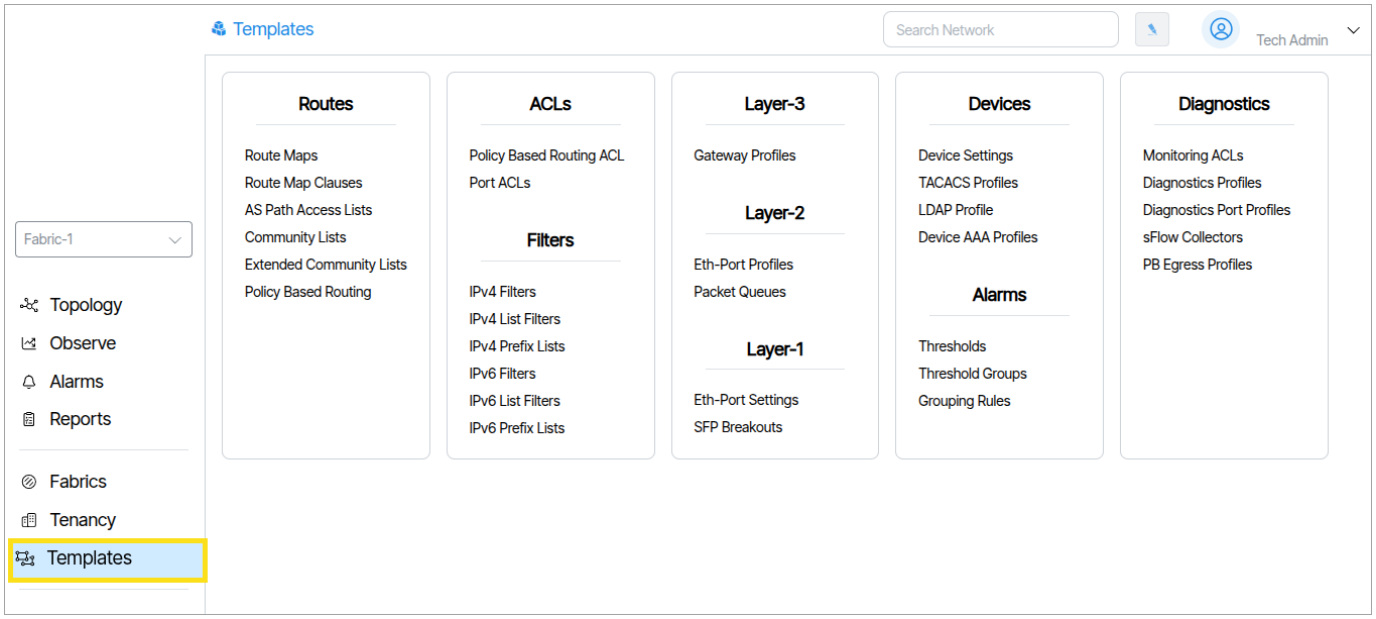


Copyright BE Networks 2026

4.5 Templates

4.5.1 Templates

The **Templates** section is composed of tools used to configure and create reusable network device settings and policies. The reusable collections are called **Template Objects**. Templates and Template Objects are organized into functional categories based on network layers and operational domains.



How to Create a Template Object

Read: Create Provisioning Objects Through Forms:

Categories

Routes

Templates for controlling routing behavior, including route maps for policy-based routing decisions, AS path filtering for BGP route selection, and community-based route tagging and filtering policies.

ACLs

Access control list templates for policy-based routing decisions and port-level traffic filtering, defining which traffic flows are permitted or denied based on specified criteria.

Filters

Network prefix and address filtering templates for both IPv4 and IPv6, used to control route advertisements, traffic matching, and network access at the IP layer.

Layer-3

Network layer templates defining gateway configurations and site-level routing policies that control how traffic is forwarded between network segments.

Layer-2

Data link layer templates for Ethernet port configurations and quality-of-service packet queue management, controlling how frames are processed and prioritized.

Copyright BE Networks 2026

Layer-1

Physical layer templates defining Ethernet port physical characteristics and SFP module breakout configurations for high-speed interfaces.

Devices

Device-level configuration templates that define system-wide settings applicable across network devices.

Alarms

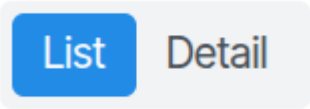
Monitoring and alerting templates that establish threshold values, organize thresholds into logical groups, and define rules for classifying alarm conditions.

Diagnostics

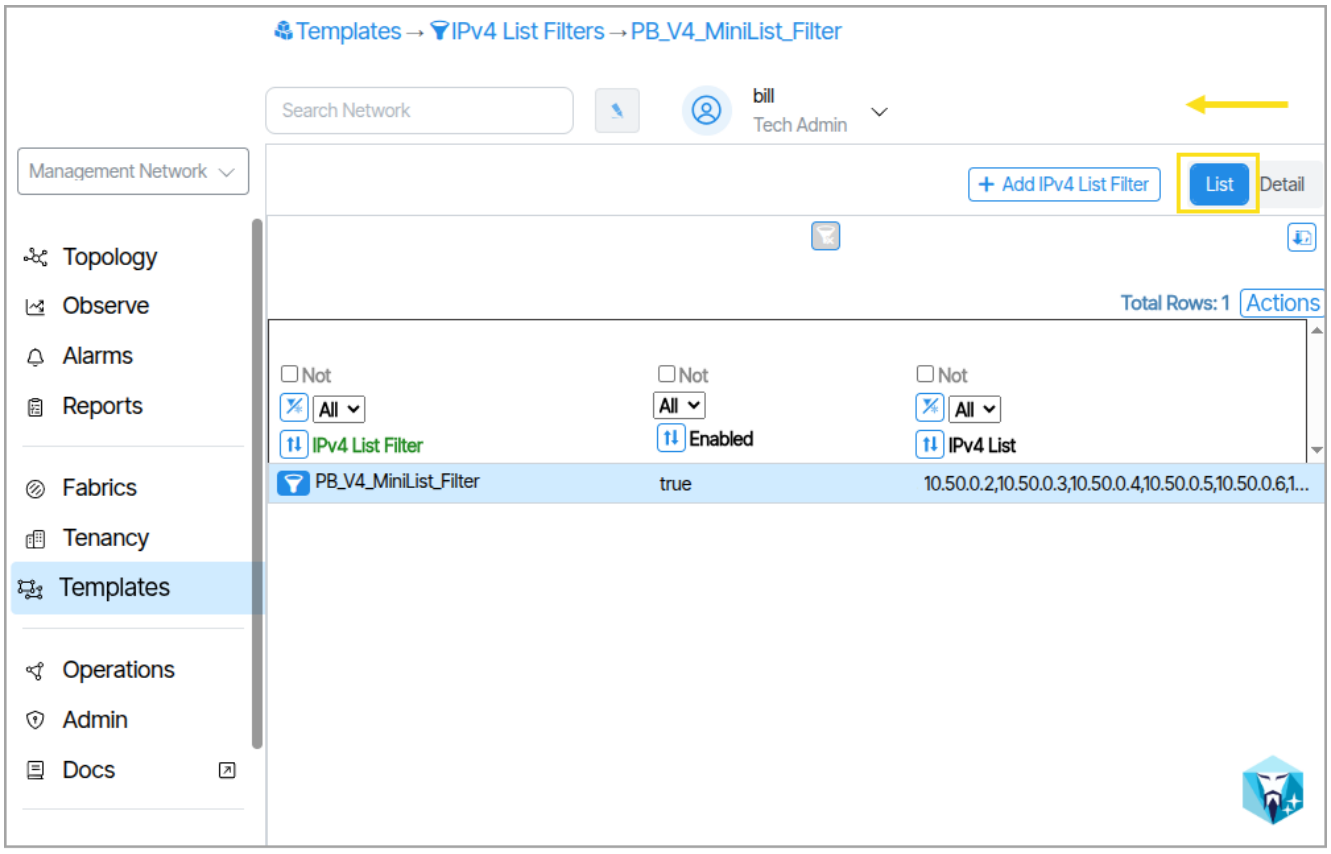
Network troubleshooting and monitoring templates for diagnostic operations, port-level diagnostics, sFlow traffic sampling configurations, and packet broker egress policies.

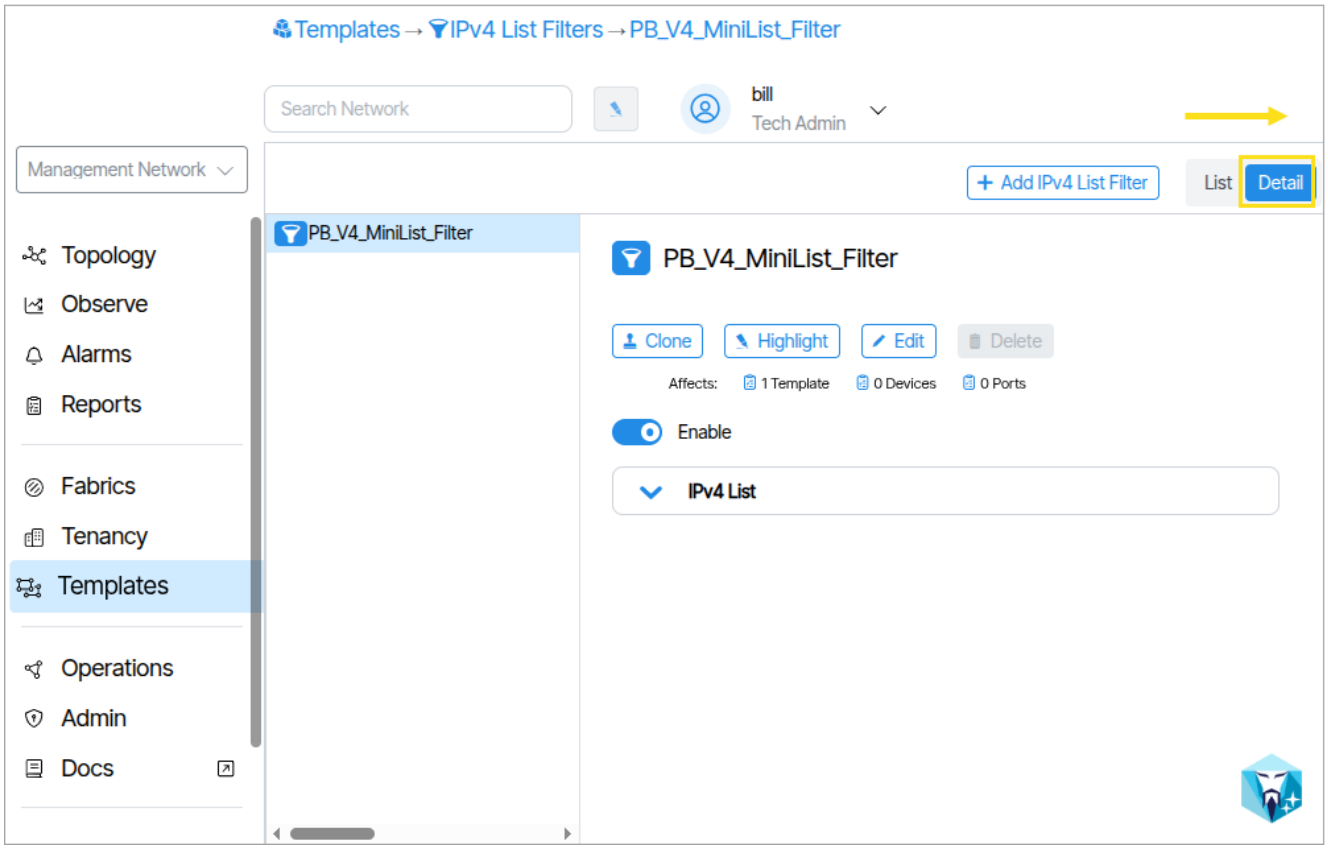
List / Detail

The List and Detail tabs display the same information in two distinct styles: **List** provides a report view, while **Detail** offers a layout better suited to direct editing.



Copyright BE Networks 2026





Copyright BE Networks 2026

ACL Routing

4.5.2 Route Maps

> **Note**
>
> Route Maps are available at Templates → Routes

A **Route Map** is a policy tool that allows administrators to override automated path selections and force specific traffic flows to follow predetermined routes. The purpose of using route maps is to direct traffic toward paths that are faster, more efficient, and experience less packet loss than the automatically selected routes.

Route-maps are used to filter or modify routes based on many possible criteria. Route-maps are composed of one or more **Route-map Clauses**. Each clause uses a filter to select the relevant routes (match criteria). Once selected, parameters of the routes can be modified with Set commands.

> **Beware of Misconfigurations**
>
> It is important that network administrators are deeply familiar with BGP and Route Map concepts before manipulating settings. Misconfiguration can lead to routing loops, traffic disruption, service outages, security vulnerabilities, and performance degradation.

Route Map Composition

Route Maps are composed of:

Route Map Clause(s) – This is a statement or rule within the route map that is composed of *permit* or *deny match conditions*, including matches to **Lists**.

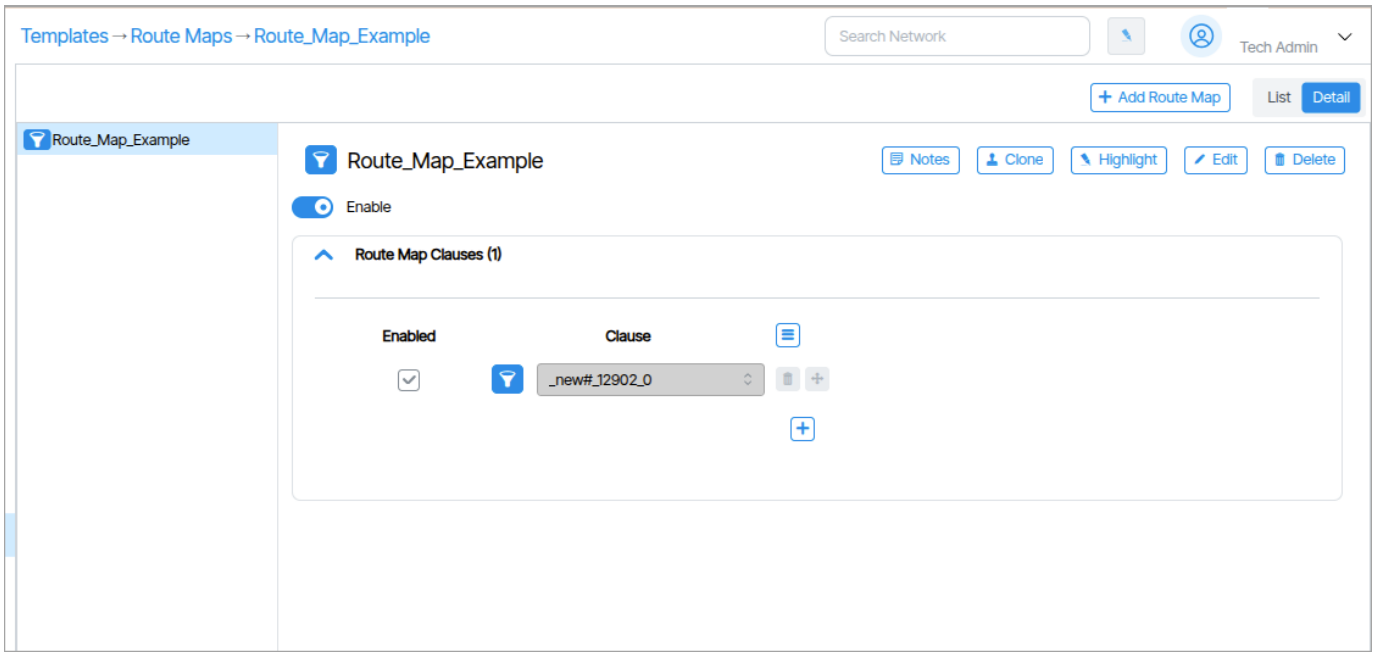
Access Control List(s) – A collection of **Conditions** or criteria that define specific attributes to match against network traffic. (ACLs are used to match 5 tuple)

Condition(s) – These are individual factors such as IP addresses, destination IP addresses or protocol types.

CREATING ROUTE MAPS

> **Creating Objects**
>
> Step-by-step instructions for creating new provisioning objects: Creating and Configuring Provisioning Objects.

Copyright BE Networks 2026



Double-click the input field that says **Choose a Route Map Clause** and select **Create New Route Map Clause** in the select box. 

Adding Multiple Clauses

To create and enable multiple Route Map Clauses, click the **Add Row** button.

Route Maps Clauses

Route Map Clauses are composed of **permit** or **deny** match conditions, including matches to **Lists** in **Match Fields**.

CREATE A ROUTE MAP CLAUSE

1. Click **Templates** and then click **Route Map Clauses**.
2. Click the **Add Route Map Clause** button.
3. In the window that appears type the name of the Route Map Clause and click **Create Route Map Clause**.

Match Conditions in a Route Map Clause

1. Click **Templates** and click **Route Map Clauses**.
2. In the fields that appears create a new Match field by clicking the **Select a Match Field to Add**.
3. In the window that appears, select a choice from the drop down menu.

Access Control Lists

An access control list (ACL) is a collection of conditions to match specific attributes of network traffic. These conditions can include source IP addresses, destination IP addresses, protocol type and other data.

CREATING AN ACCESS CONTROL LIST

AS Path Lists

AS Path Lists allow a route map to select routes with a specific AS number in the path using regex.

Copyright BE Networks 2026

CREATING AN AS PATH LIST

1. Click **Templates** and click **AS Path Access List**.
2. Click the **Add AS Path Access List** button.
3. Enter a name for the list and click **Create ne As Path Access List**.
4. Set the conditions for the list.
5. In the **Edit Mode** prompt click the **Save** button.

Community Lists

Community Lists allow a route map to select routes originating from a specific community.

CREATING A COMMUNITY LIST

1. Click **Templates** and click **Community List**.
2. Click the **Add a New Community List** button.
3. Enter a name for the list and click **Create Community List**.
4. Set the conditions for the list.
5. In the **Edit Mode** prompt click the **Save** button.

Extended Community Lists

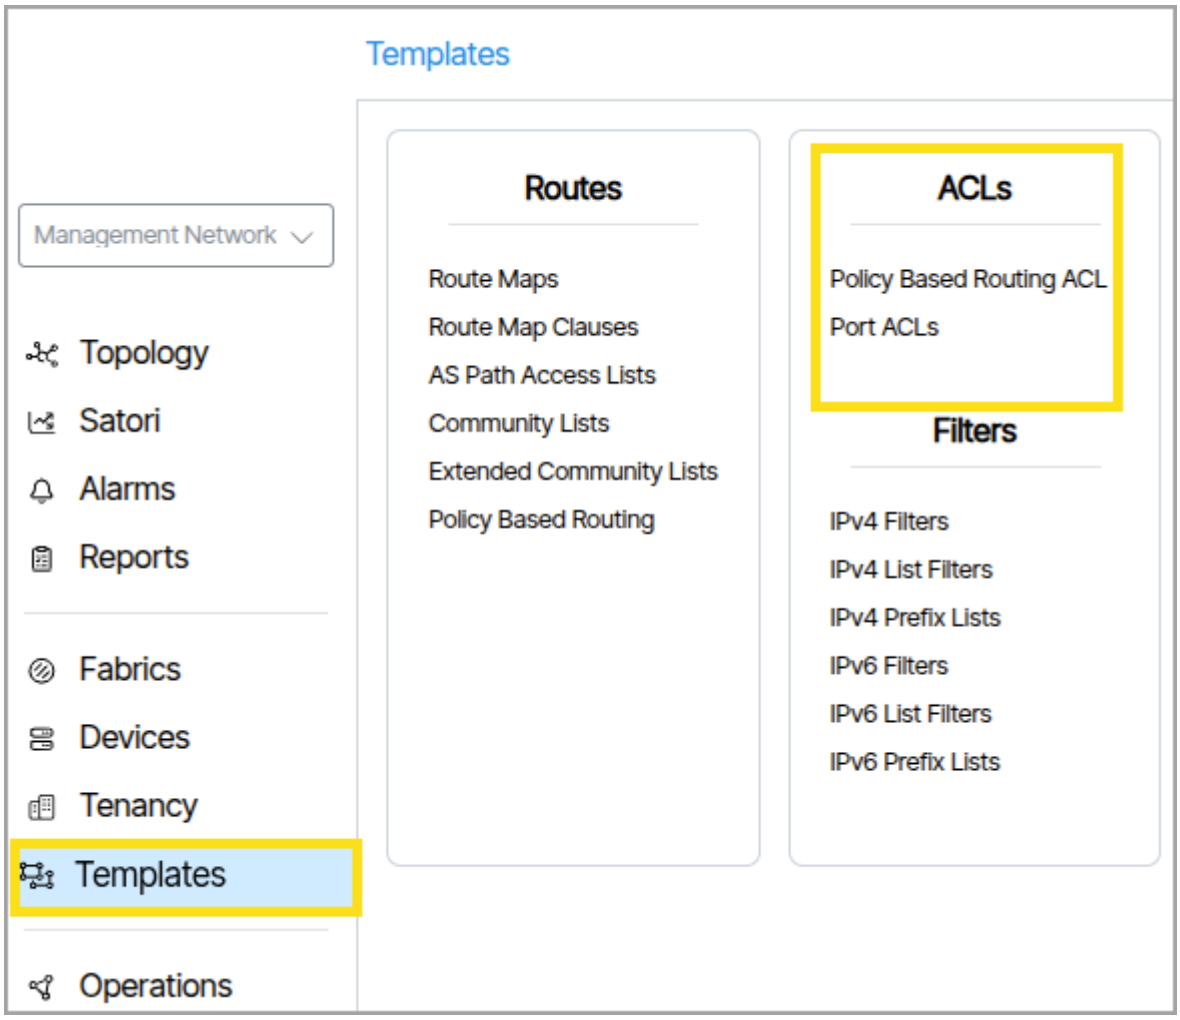
An Extended Community List is a filter that controls how routes are shared between separate virtual networks (VRFs) on a router.

CREATING AN EXTENDED COMMUNITY LIST

1. Click **Templates** and click **Extended Community List**.
2. Click the **Add a New Extended Community List** button.
3. Enter a name for the list and click **Create Extended Community List**.
4. Set the conditions for the list.
5. In the **Edit Mode** prompt click the **Save** button.

Copyright BE Networks 2026

4.5.3 ACLs



An **ACL (Access Control List)** is a collection of IP address filters. Verity lets administrators configure ACLs at the Service and Port level. Applying an ACL to a Service and/or Port is composed of these steps:

- Create and enable a Port ACL.
- Create and configure IPv4 and/or IPv6 filters.
- Apply filters to the Port ACL.
- From within a chosen Eth-Port profile, assign the ACL to a Port and/or Service(s).

Policy Based Routing ACL

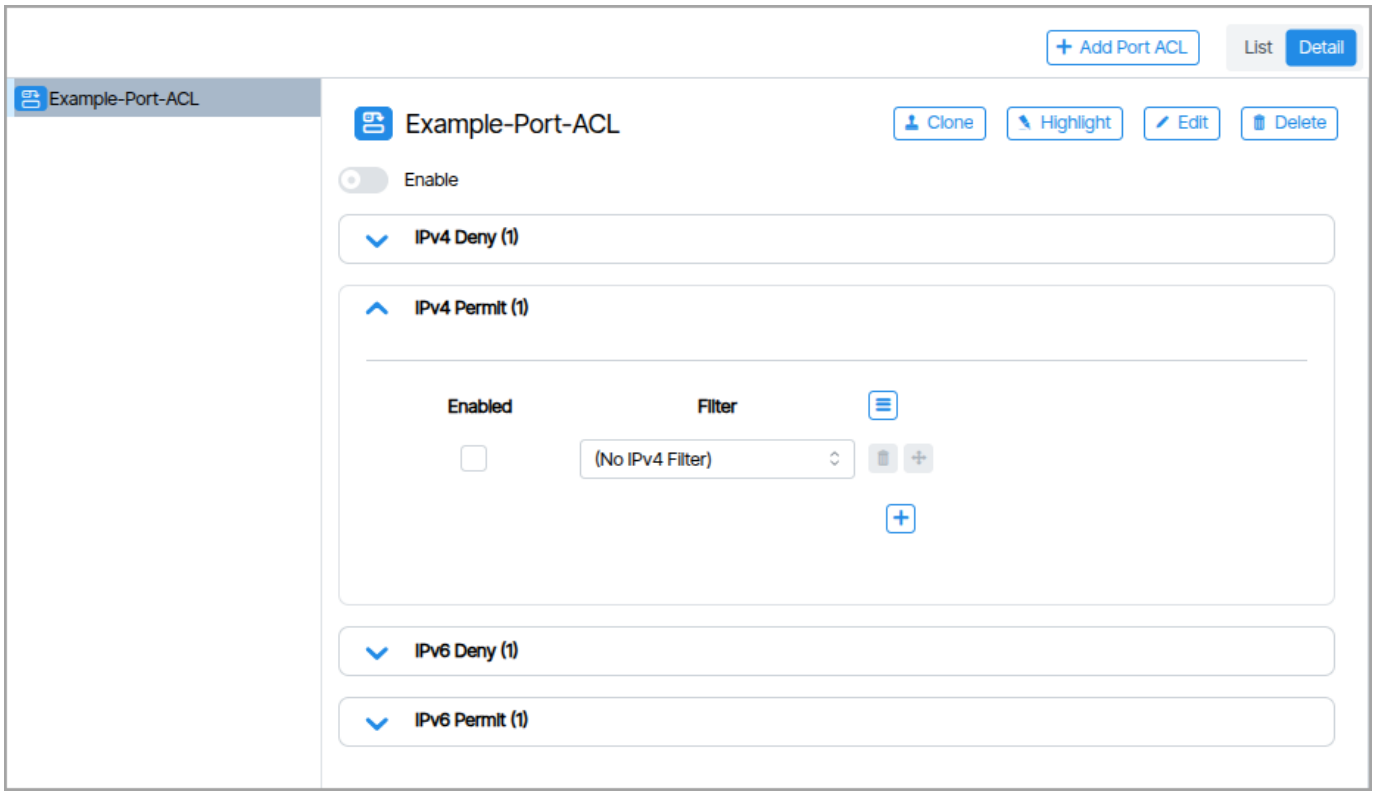
Port ACLs

CREATE PORT ACLS

1. Double-click **Templates** and click **Port ACLs**.
2. Click the **Add a New Port ACL** button.
3. Give the Port ACL a name and click **Create Port ACL**.

The window that appears is composed of sections titled IPv4 Deny, IPv4 Permit, IPv6 Deny, and IPv6 Permit. These sections contain filters that determine whether selected network data is forwarded to a destination device or blocked.

Copyright BE Networks 2026

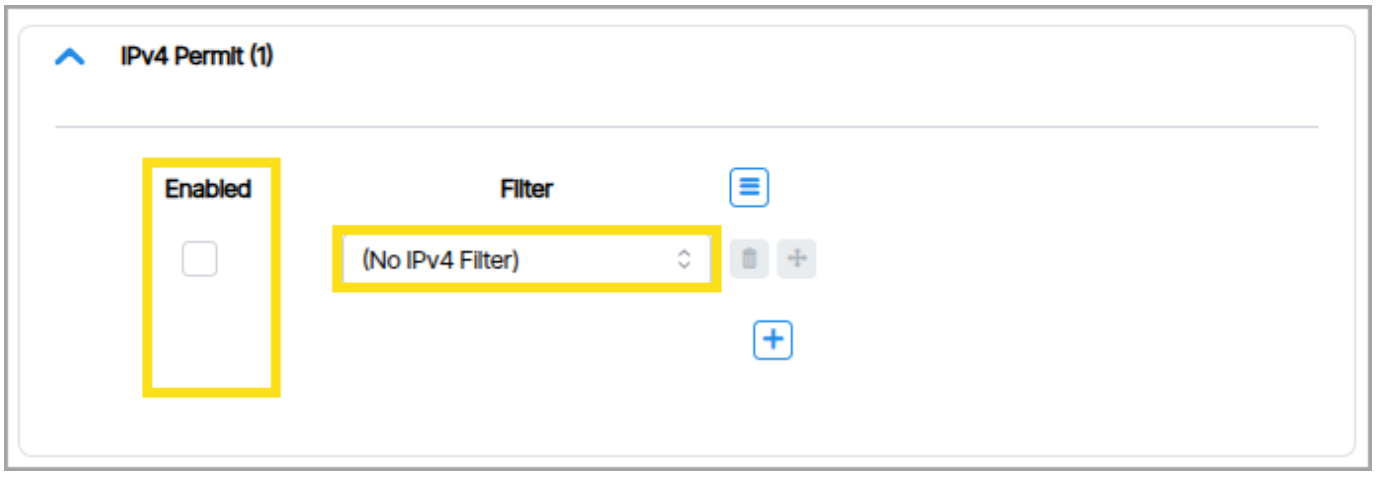


Create and Configure Filters

Before applying filters to an ACL you must first create and configure the filters. How to Create IPV4 and IPV6 Filters

Apply Filters to the Port ACL

In the Port ACL you apply the filters by setting them under the field named **Filter**. You activate the Filter by checking the **Enabled** box.

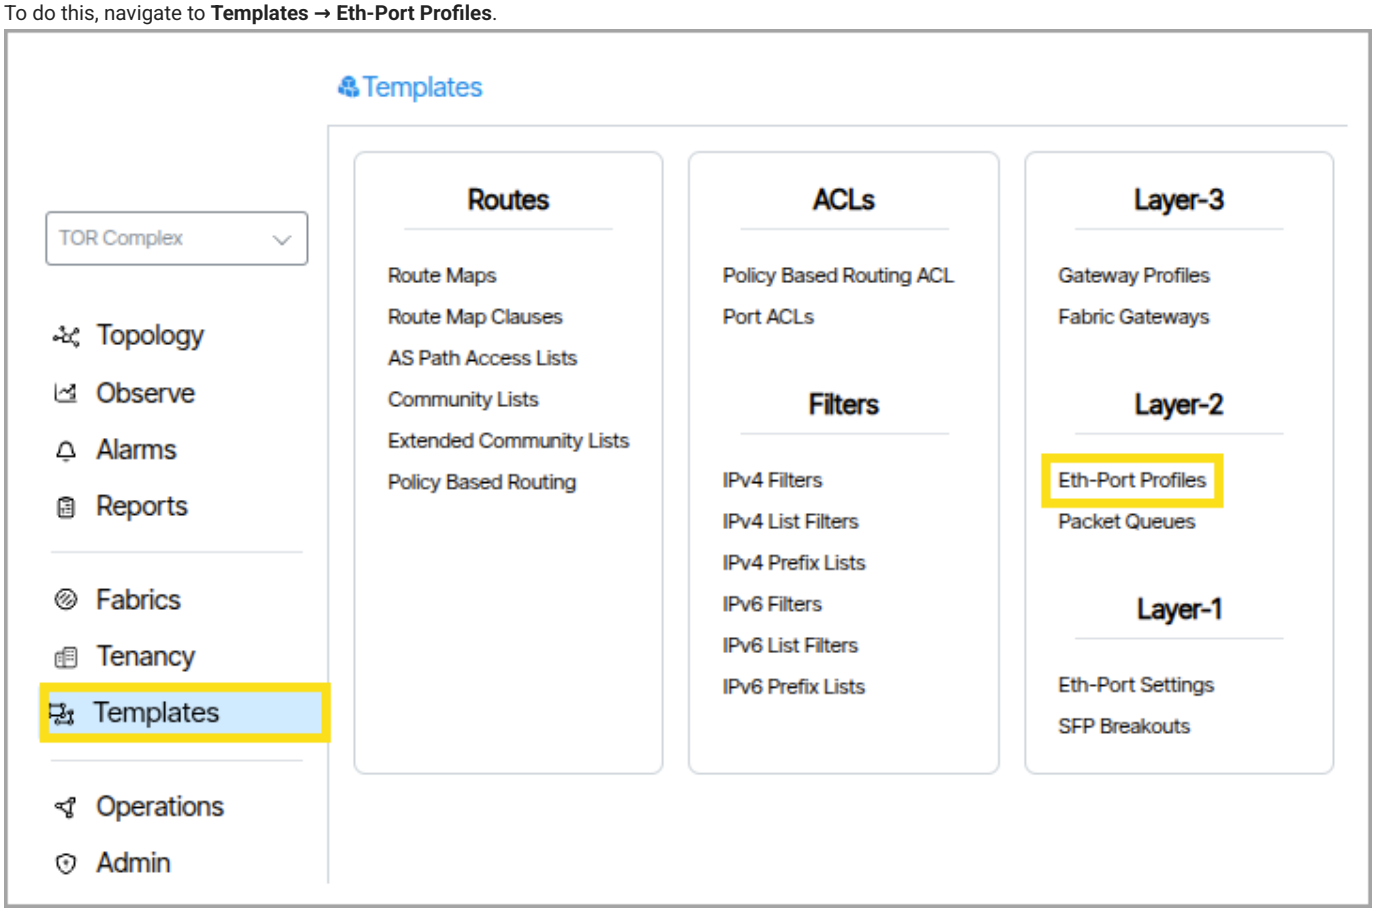


How to Assign a Port ACL to a Port or Service

After assigning filters to Port ACL(s), you apply the Port ACL to the Service and/or Port within a selected Eth-Port Profile.

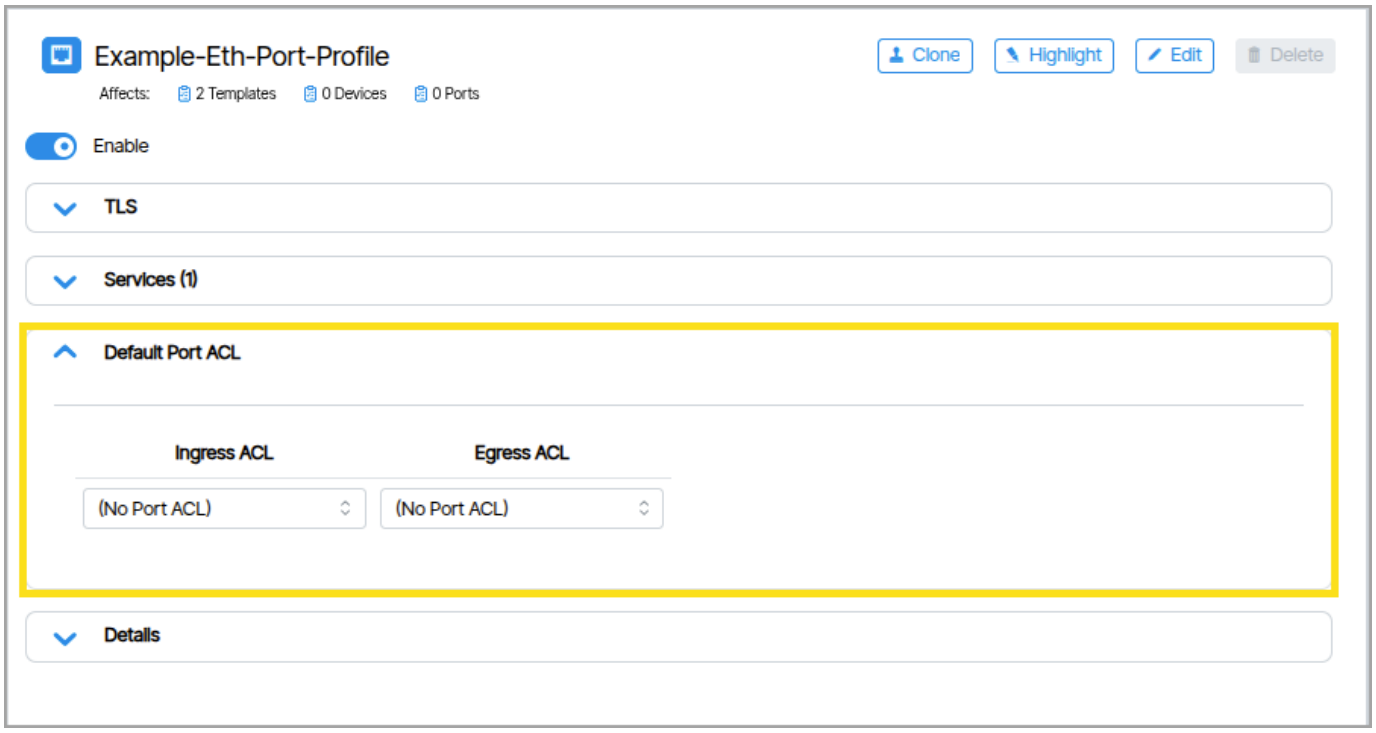
Copyright BE Networks 2026

To do this, navigate to **Templates → Eth-Port Profiles**.



ASSIGNING ACL TO A PORT

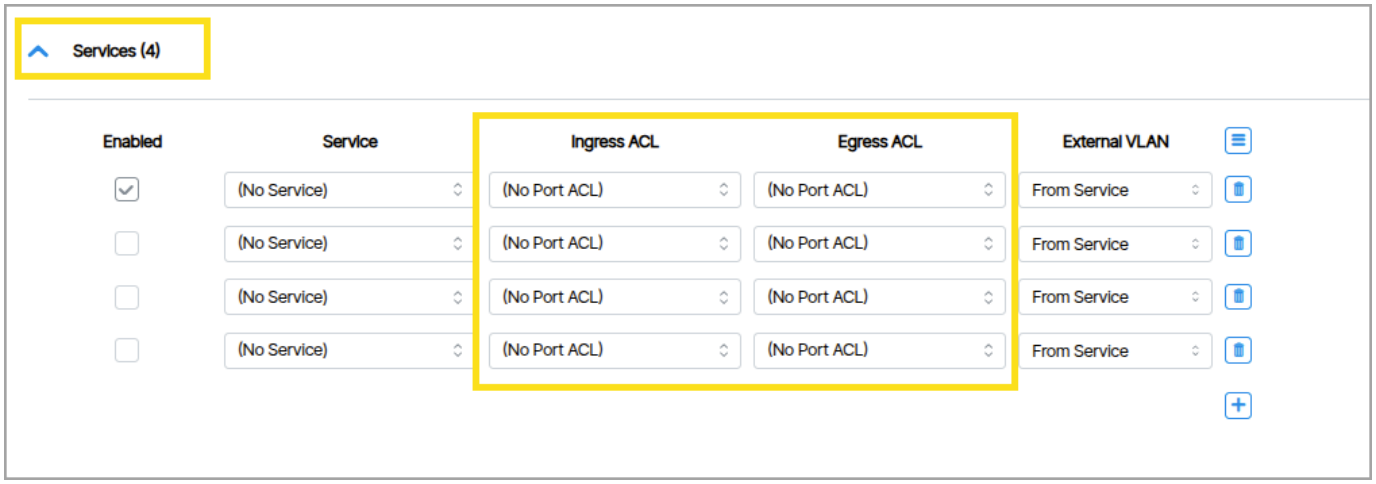
Assigning ACL to the Port is performed in the **Port ACLs Ingress ACL** and **Egress ACL** settings.



Copyright BE Networks 2026

ASSIGNING ACL TO A SERVICE

Each service in the **Eth-Port Profile** has Ingress and Egress ACL fields. These fields are where you set the ACL Port.



Common processing order

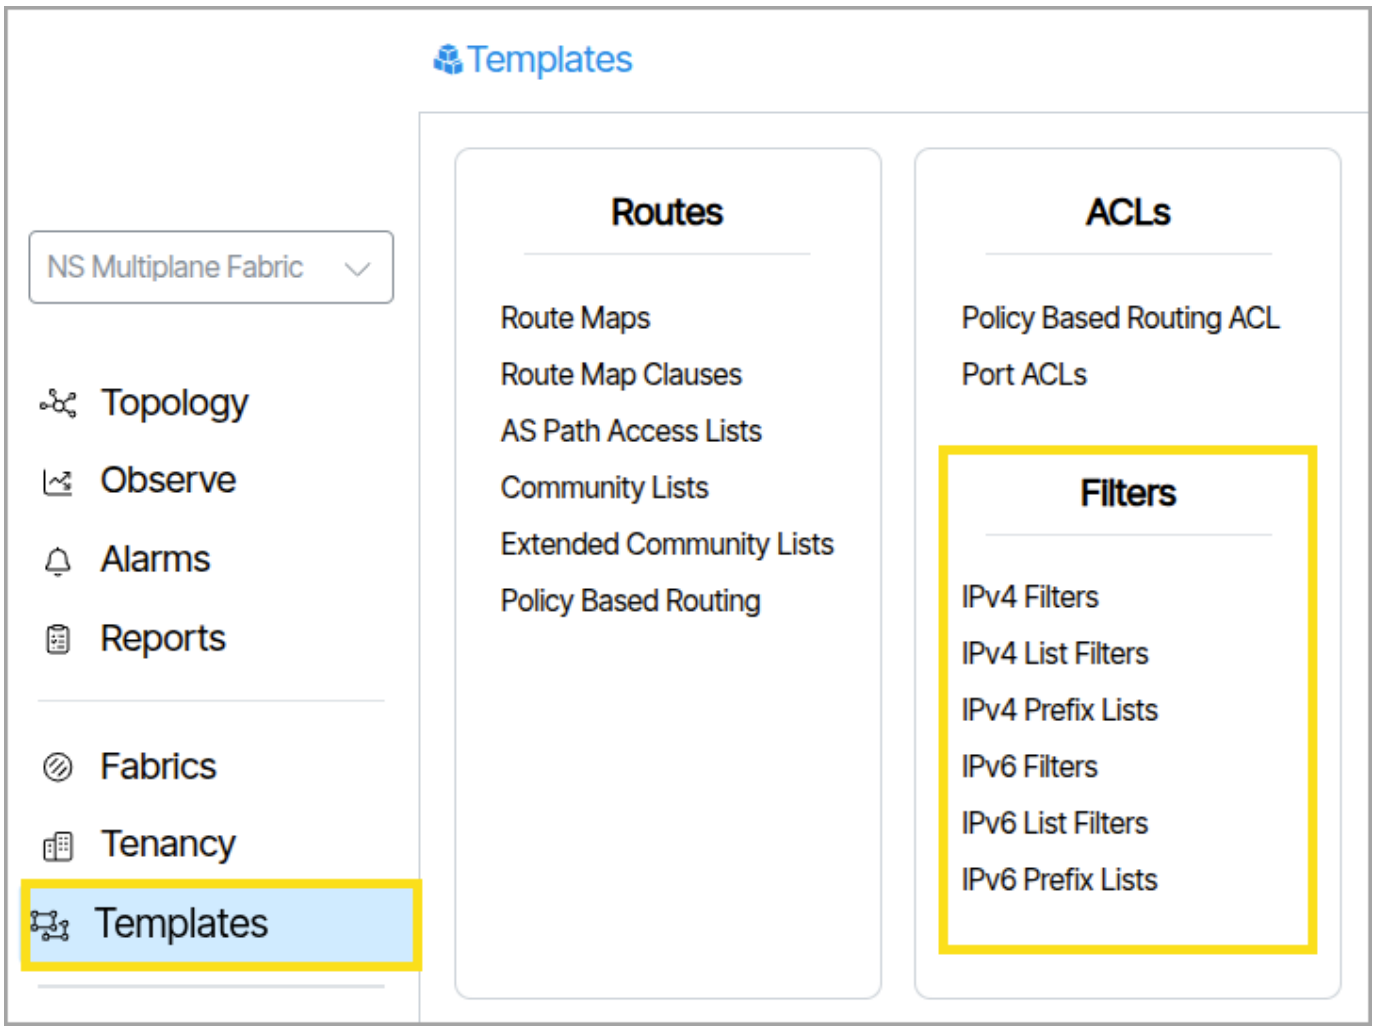
1. **Port ACL** is applied first (affects all VLANs on the port)
2. **Service ACL** is applied second (affects only the specific VLAN)
3. Traffic must satisfy both ACLs to be permitted

Copyright BE Networks 2026

Filter

4.5.4 Filters

This section provides tools that enable administrators to create IP filters controlling network traffic and routing information based on IP address matching. Administrators configure these filters to permit, deny, or process traffic as required, then apply them to other objects throughout the application.



Filter Types

- IPv4 Prefix Lists
- IPv6 Prefix Lists
- IPv4 List Filters
- IPv4 Filters
- IPv6 List Filters
- IPv6 Filters

Copyright BE Networks 2026

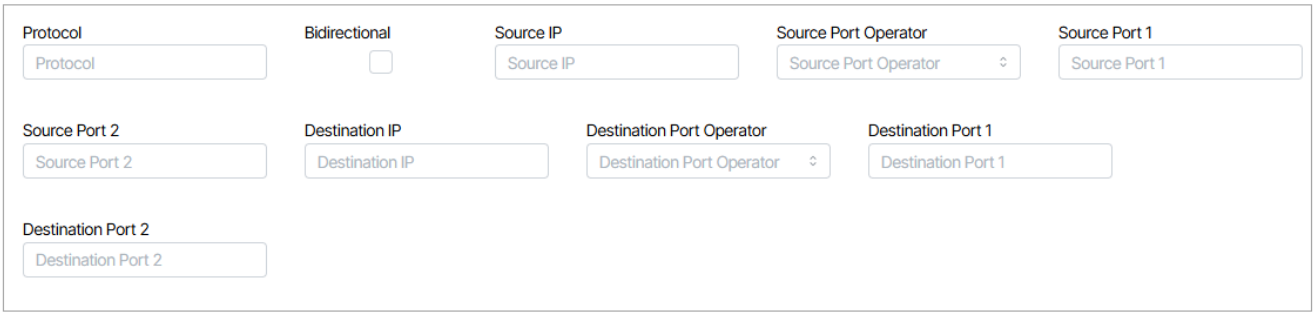
CREATE FILTERS

To create any filter type in the Filters list you follow this process:

1. Double-click **Templates** and click the desired filter type ().
2. Click the **Add** button .
3. Name the filter and click the **Create** button .
4. In the filter settings, edit the rules and click the **Enabled** box to enable the filter ().

IPv4/IPv6 Filters

Filter Rules

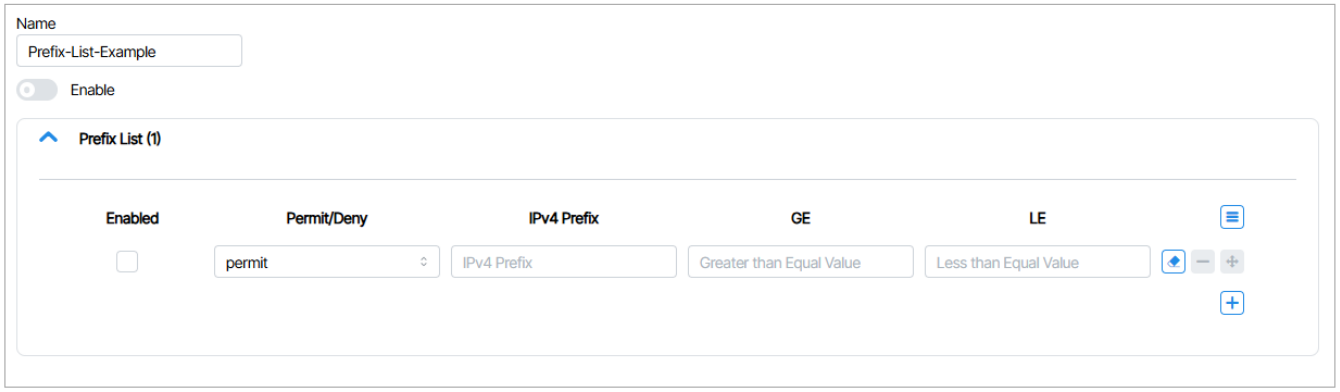


Filter	Function
Protocol	IP protocol by name (ip, tcp, udp, or icmp, only) or IP Protocol by IANI number (0-255).
Bidirectional	Bidirectional packets if selected, else unidirectional.
Source IP	IPv4 or IPv6 packet source address. Use IP mask for range of IPs.
Source Port Operator	This field determines which match operation will be applied to TCP/UDP port(s). The choices are equal, greater-than, less-than or range.
Source Port 1	This field is used for equal, greater-than or less-than TCP/UDP port value in match operation. This field is also used for the lower value in the range port match operation.
Source Port 2	This field will only be used in the range TCP/UDP port value operation to define the top value in the range
Destination IP	IPv4 or IPv6 packet source address. Use IP mask for range of IPs.
Destination Port Operator	This field determines which match operation will be applied to TCP/UDP ports. The choices are equal, greater-than, less-than or range.
Destination Port 1	This field is used for equal, greater-than or less-than TCP/UDP port value in match operation. This field is also used for the lower value in the range port match operation.
Destination Port 2	his field will only be used in the range TCP/UDP port value match operation to define the top value in the range.

IPv4/IPv6 Prefix Lists

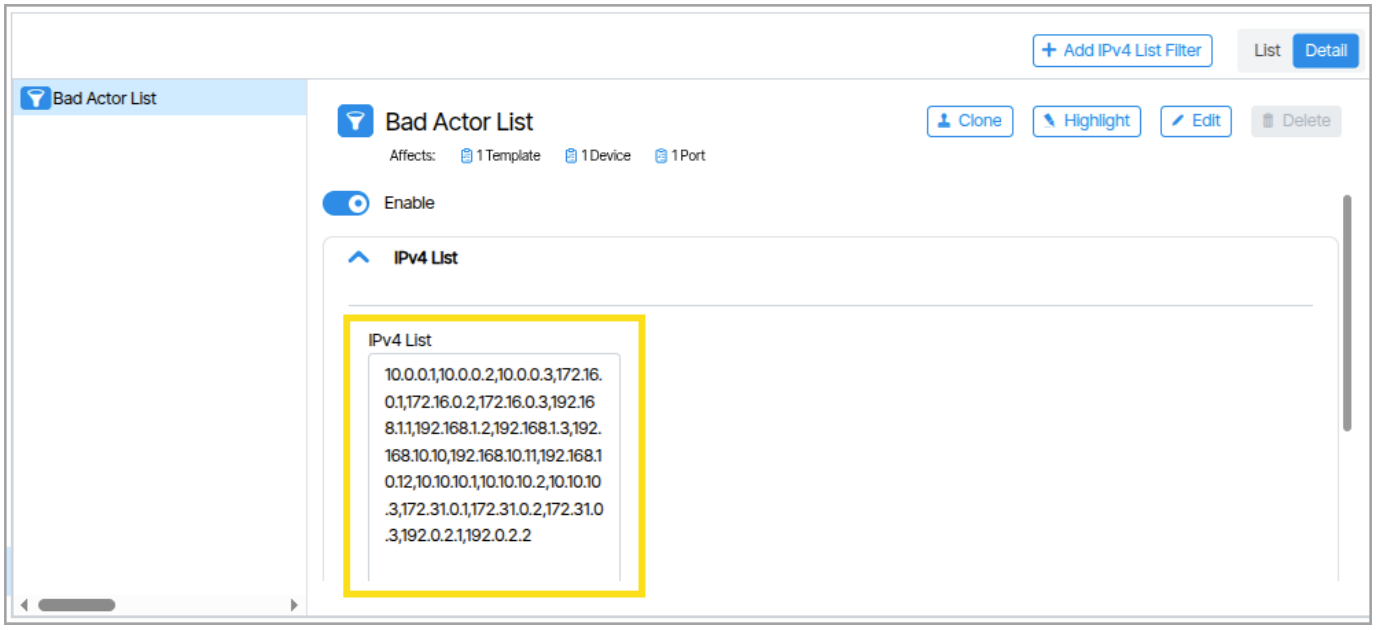
A Prefix List is a structured list of IPv4 or IPv6 network prefixes that is used to filter or match the IPv4 or IPv6 addresses.

Copyright BE Networks 2026



IPv4/IPv6 List Filters

To filter individual IP addresses or lists of individual IP addresses, use an **IPv4 List Filter** or **IPv6 List Filter** and enter the addresses in the provided field. A comma must be used between each IP address. There are no spaces used in the lists.

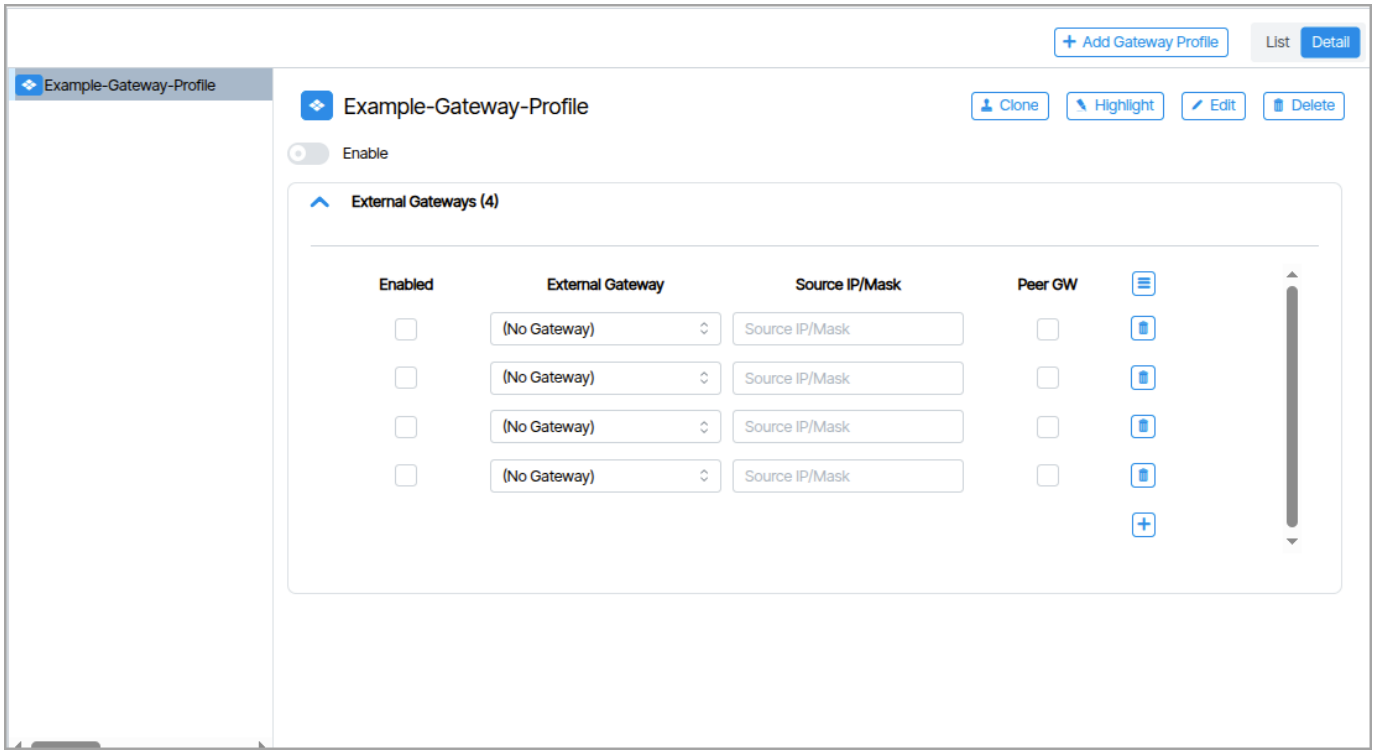


Copyright BE Networks 2026

4.5.5 Layer-3

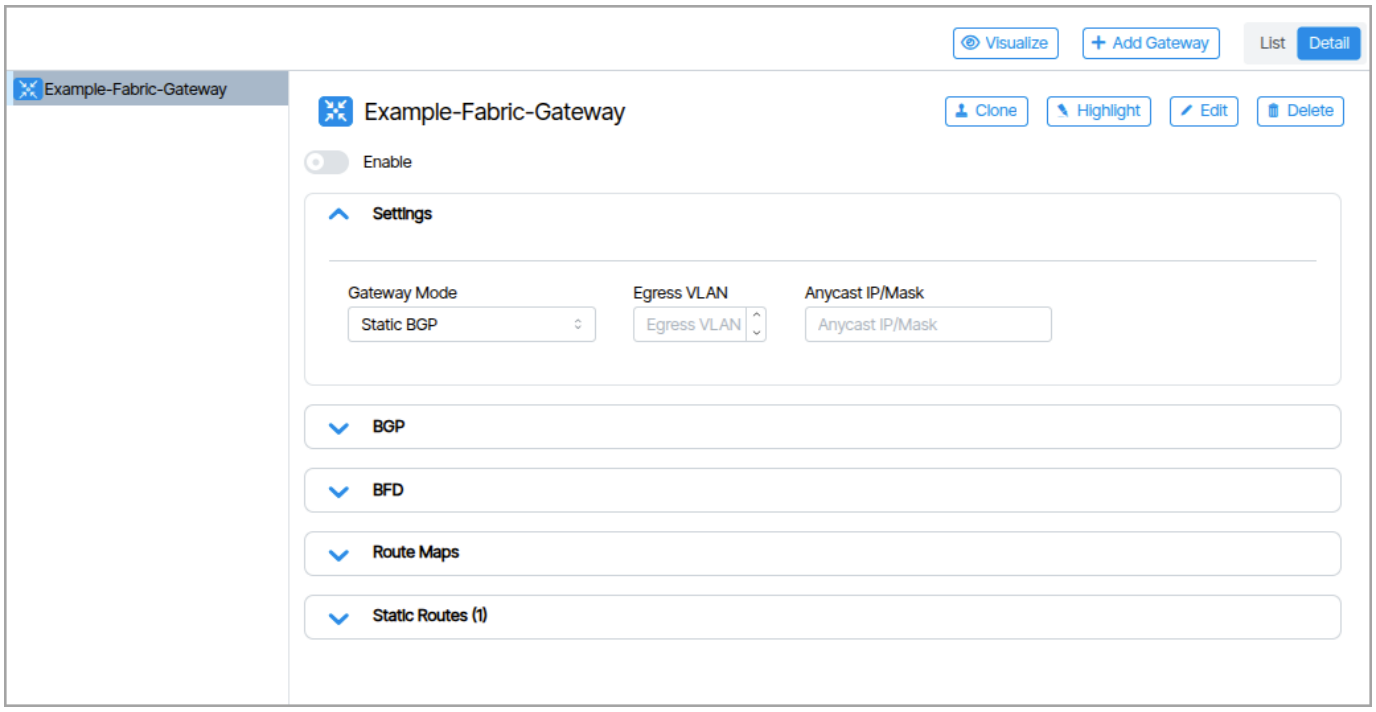
Gateway Profiles

A BGP Gateway Profile is a packaging of one or more BGP Tenant Gateways that are served by this port.



Fabric Gateways

A Fabric Gateway is a configuration object that contains settings for traffic between a local site (like a branch office, campus, or data center) and external networks.



Copyright BE Networks 2026

4.5.6 Related Information

Gateways

Inter-Site Routing

Copyright BE Networks 2026

Service VLAN

4.5.7 Layer-2

Eth-Port Profiles

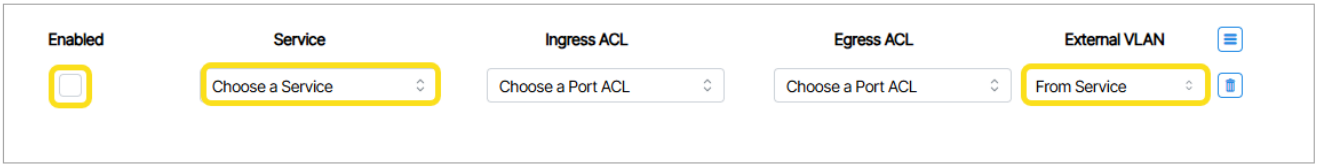
An **Eth-Port Profile** (Ethernet Port Profile) is a reusable configuration template that packages multiple services for deployment on a port. This feature simplifies port provisioning by allowing operators to group related Layer-2 VLAN services commonly required by servers and apply them as a single unit.

Creating an Eth-Port Profile

1. Navigate to the Eth-Port Profiles section by double-clicking **Templates** and clicking **Eth-Port Profiles**.
2. Create a new Eth-Port Profile object by clicking the **Add** button.
3. In the prompt that appears, name the Eth-Port Profile and assign it a group.
4. Check the **Enable** box.

ASSIGNING SERVICES TO AN ETH-PORT PROFILE

1. From within Templates open the Eth-Port Profile tile. Select the desired Eth-Port-Profile and set the **Service** and **External VLAN** fields.
2. Enable your selection(s) by clicking the **Enabled** checkbox.

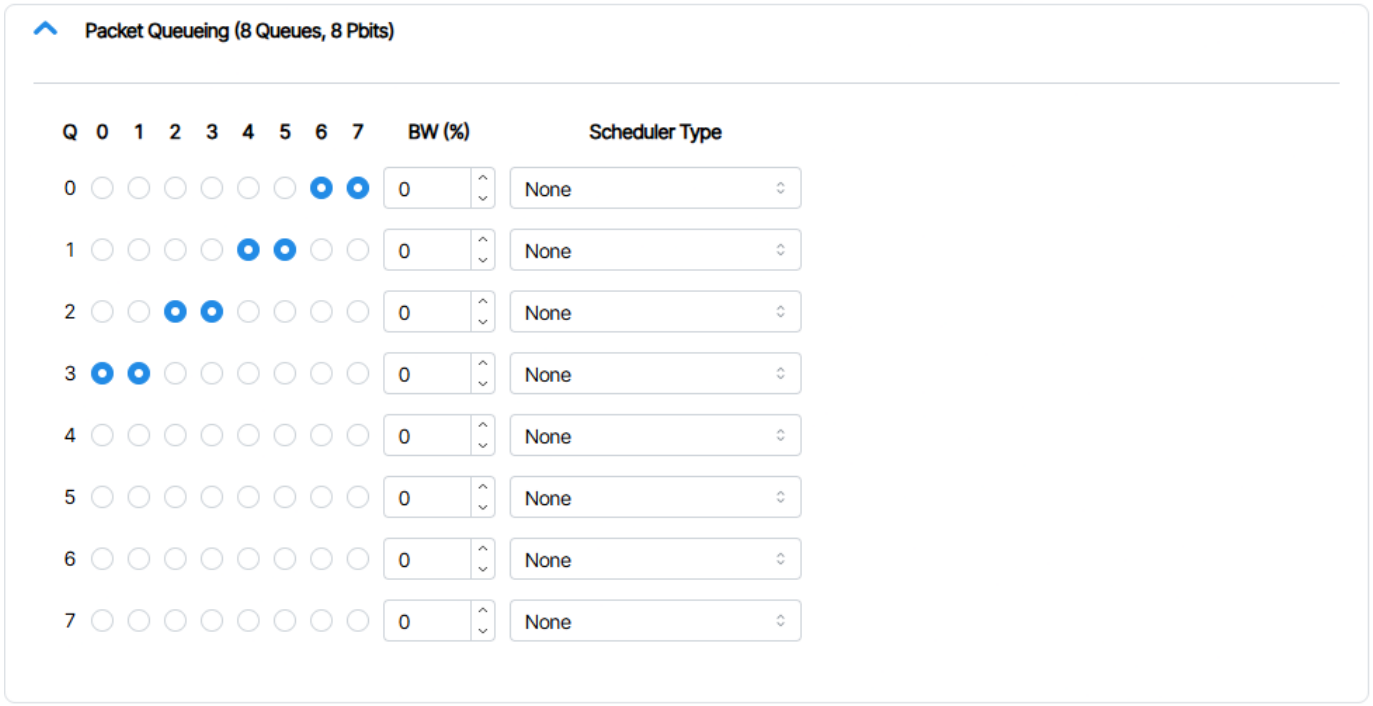


Packet Queues

This Template lets you set the priority for packets depending on bandwidth.

Here you set the priority for p-bits and set the bandwidth (BW) for each queue as a percentage. The percentage is the maximum average percentage of the port output bandwidth that packets in that queue will be allowed to consume. The value cannot exceed 100 and a value of 0 means no limit is set.

Copyright BE Networks 2026



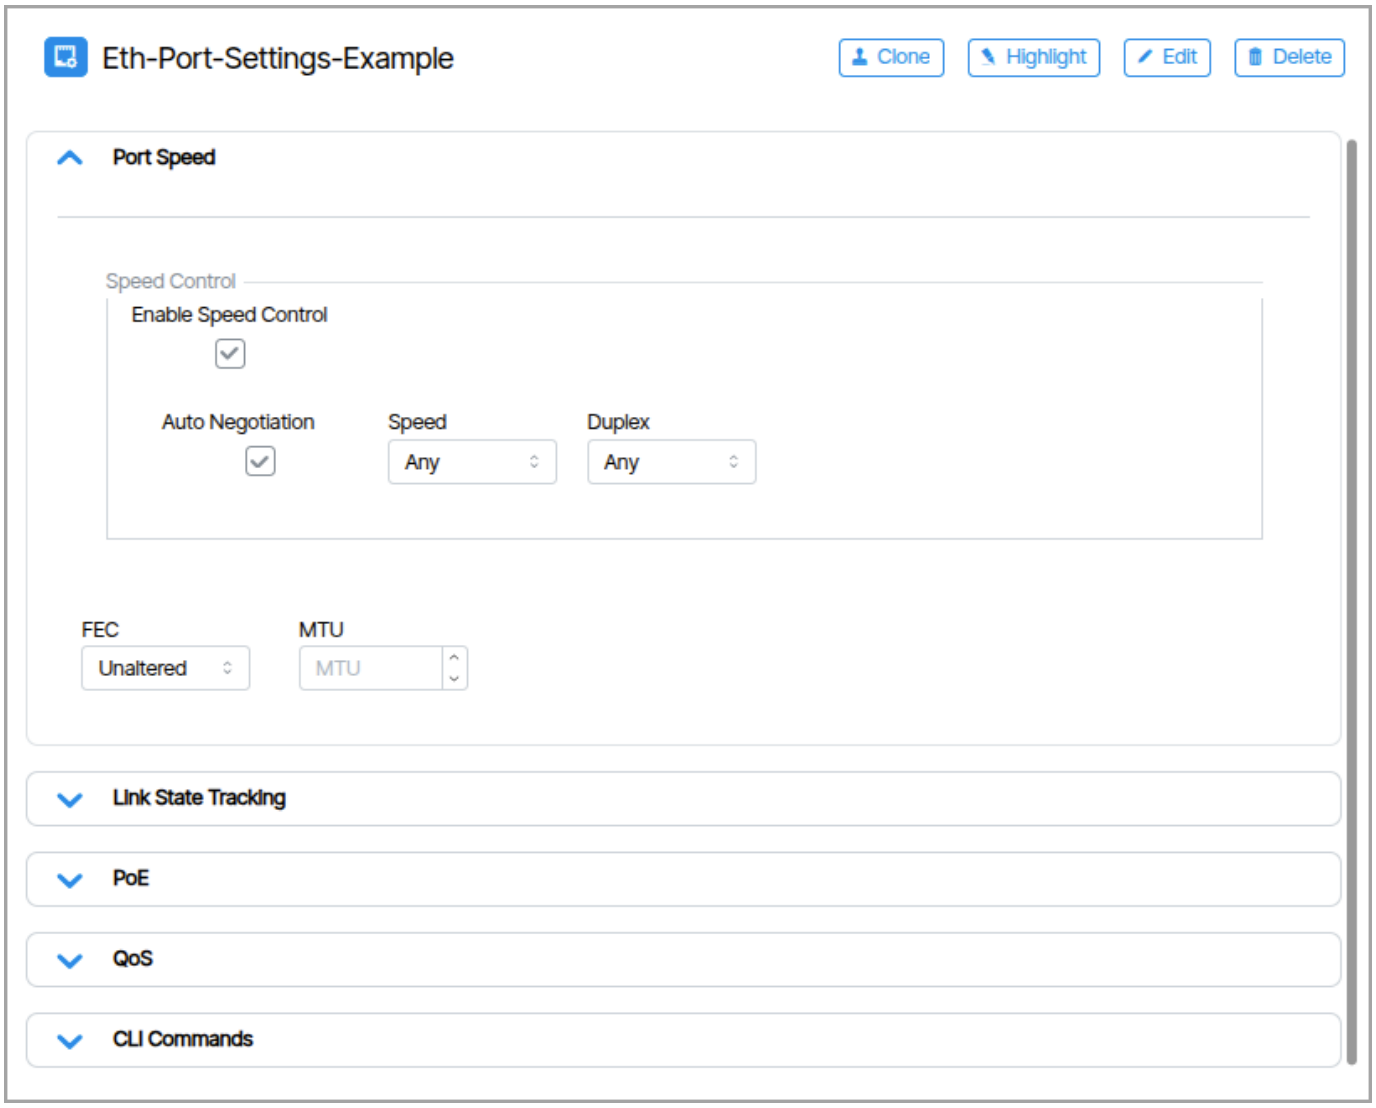
Copyright BE Networks 2026

4.5.8 Layer-1

Eth-Port Settings

Eth-Port Settings let the operator to define settings specific to a port related to Layer-1 and Layer-2 features such as port speed, speed/duplex negotiation, and Power-over-Ethernet (PoE).

The following figure shows the detail within an Ethernet Port Settings dialog box:



CREATE AN ETH-PORT SETTING

APPLY AN ETH-PORT SETTING

1. Choose the port to apply the Eth-Port Setting on ( ).
2. Open the port editing window and assign the Eth-Port Setting in the right most dropdown menu ( ).

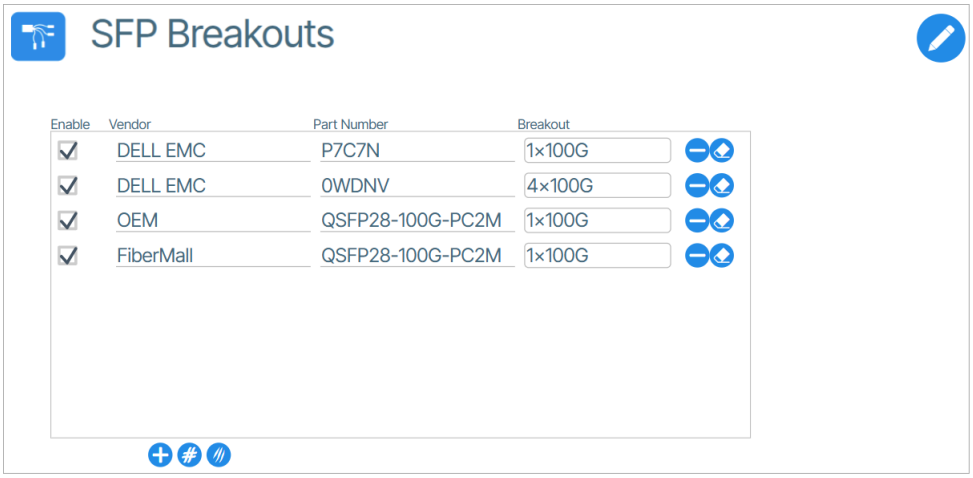
Copyright BE Networks 2026

Breakouts

BREAKOUT PORT

A **breakout port** refers to a high-density network port that can be split or "broken out" into multiple lower-speed connections using specialized cables or modules ( ).

SFP BREAKOUTS



SFP stands for Small Form-factor Pluggable. These are compact, hot-pluggable network transceivers used for both telecommunications and data communications applications.

An **SFP Breakout** refers to using a breakout cable or adapter to split a single high-speed port into multiple lower-speed SFP connections.

In Verity, the **SFP Breakouts** feature ( ) allows users to configure how SFP (Small Form-factor Pluggable) modules function within a switch. This feature can automatically detect the breakout configuration in two ways:

1. **Automatic detection**: The software reads the SFP module's specifications after it's inserted into the switch
2. **Manual configuration**: Users can define the behavior based on the module's vendor and part number

Users can configure SFP breakout behavior through the interface using either:

- **Global settings**: Apply the same configuration across all ports
- **Individual port settings**: Customize behavior for specific ports

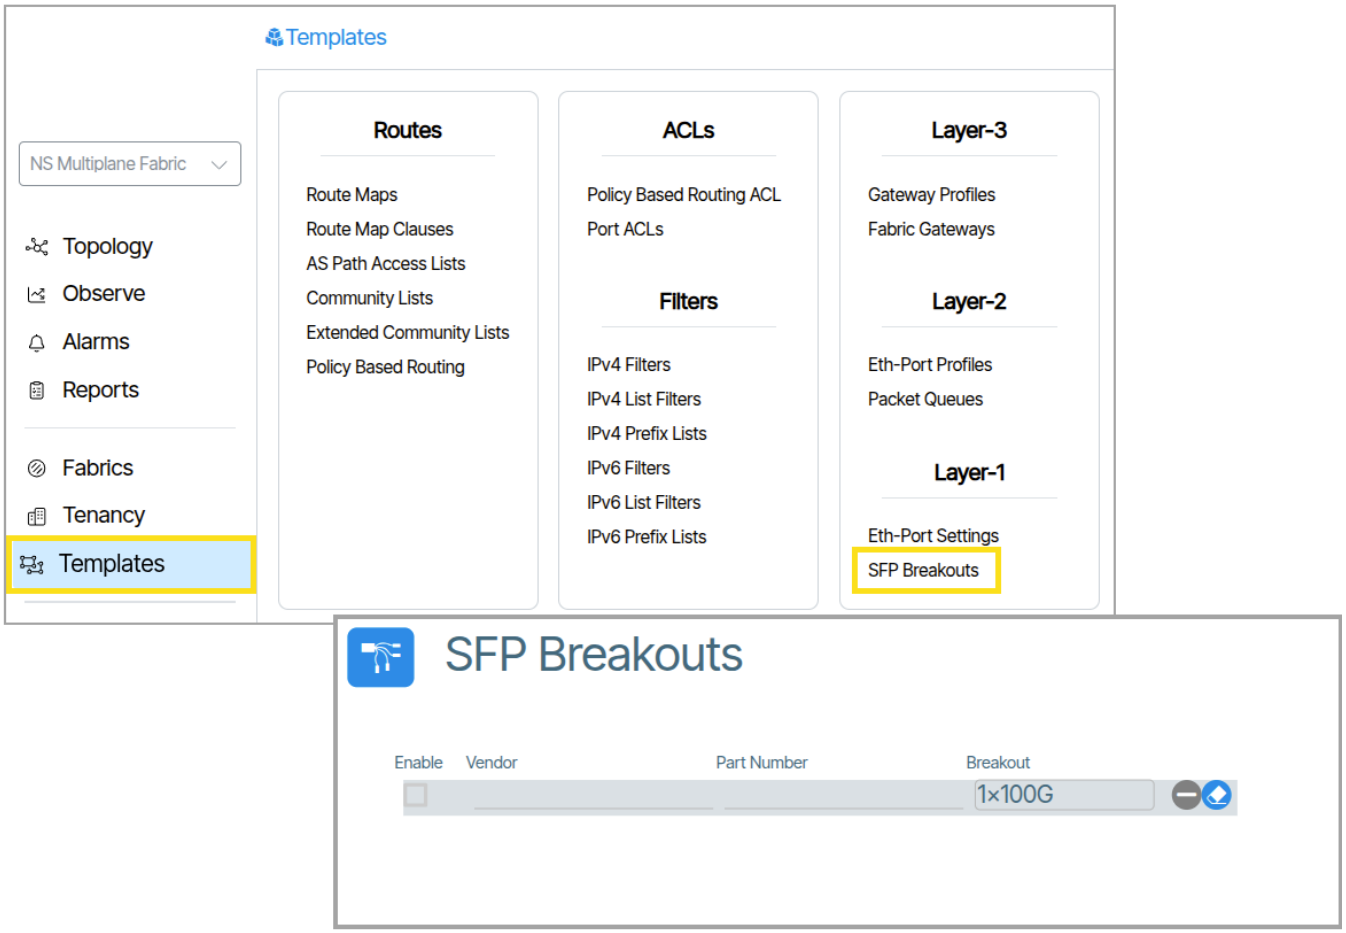
> **Warning**
>
> Breakout settings are constrained by the switch's capabilities. If the physical limitations of the switch are exceeded, an error is displayed, and the device turns red. Hovering over the device will show the following error message:
>
> **Attempting to add [*number of breakouts*] to pipeline [*pipeline index*] with [*available space*] available**.
>
> - **number of breakouts** is the count of attempted breakout additions.
> - **pipeline Index** indicates the specific pipeline number or identifier.
> - **available space** is the remaining breakouts that can be added to the pipeline.

GLOBALLY

If your intent is to set SFP device behaviors globally, throughout your system, you set this information in the Global / SFP Breakouts list. Here, you list devices by vendor and part number, and you set the behavior for each device under the field labeled *breakout*. Some SFP breakout cables identify themsleves with the same part number on each end which can result in misconfiguration based on using system wide defaults. In these cases, the

Copyright BE Networks 2026

user should set the behavior in the global section for the broken out side of the breakout cable, and use the port specific override method for the "1x" side of the cable.

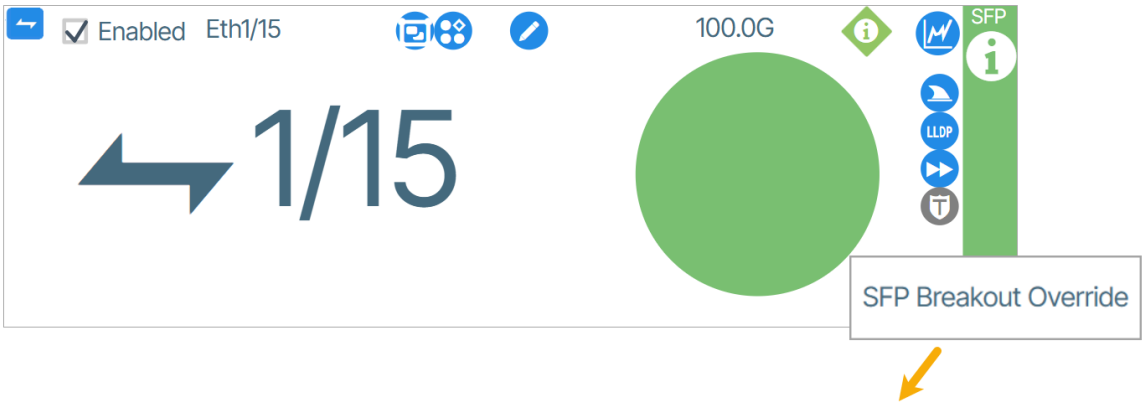


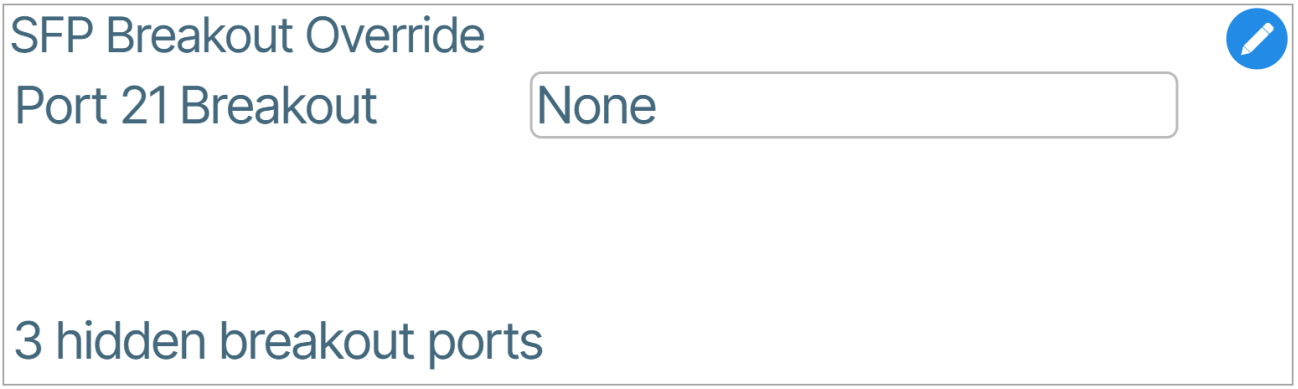
Port Specific Override

If your intent is to set an SFP device behavior on a specific port and override existing behavior, you can do so with the SFP Breakout Override form. Setting SFP devices in this way overrides existing automatic behaviors as well as the behaviors dictated in the Global Breakout settings. The drop down list in the UI is provided after the switching device has been discovered and is under management of Verity. The switch is interrogated by Verity and then The UI provides a specific drop down list based on the SFP and port it is plugged into. This precludes the ability to pre-provision port specific breakout.

It should be noted that there are limitations to the combinations and locations of multiple SFPs and the vendor's hardware specifications should be referred to in these cases. There are cases where the request from the Verity system to the switch can fail based on these limitation, and is indicated as such in the system displaying a "red fill" on the device after trying to set the port specific breakout.

Copyright BE Networks 2026





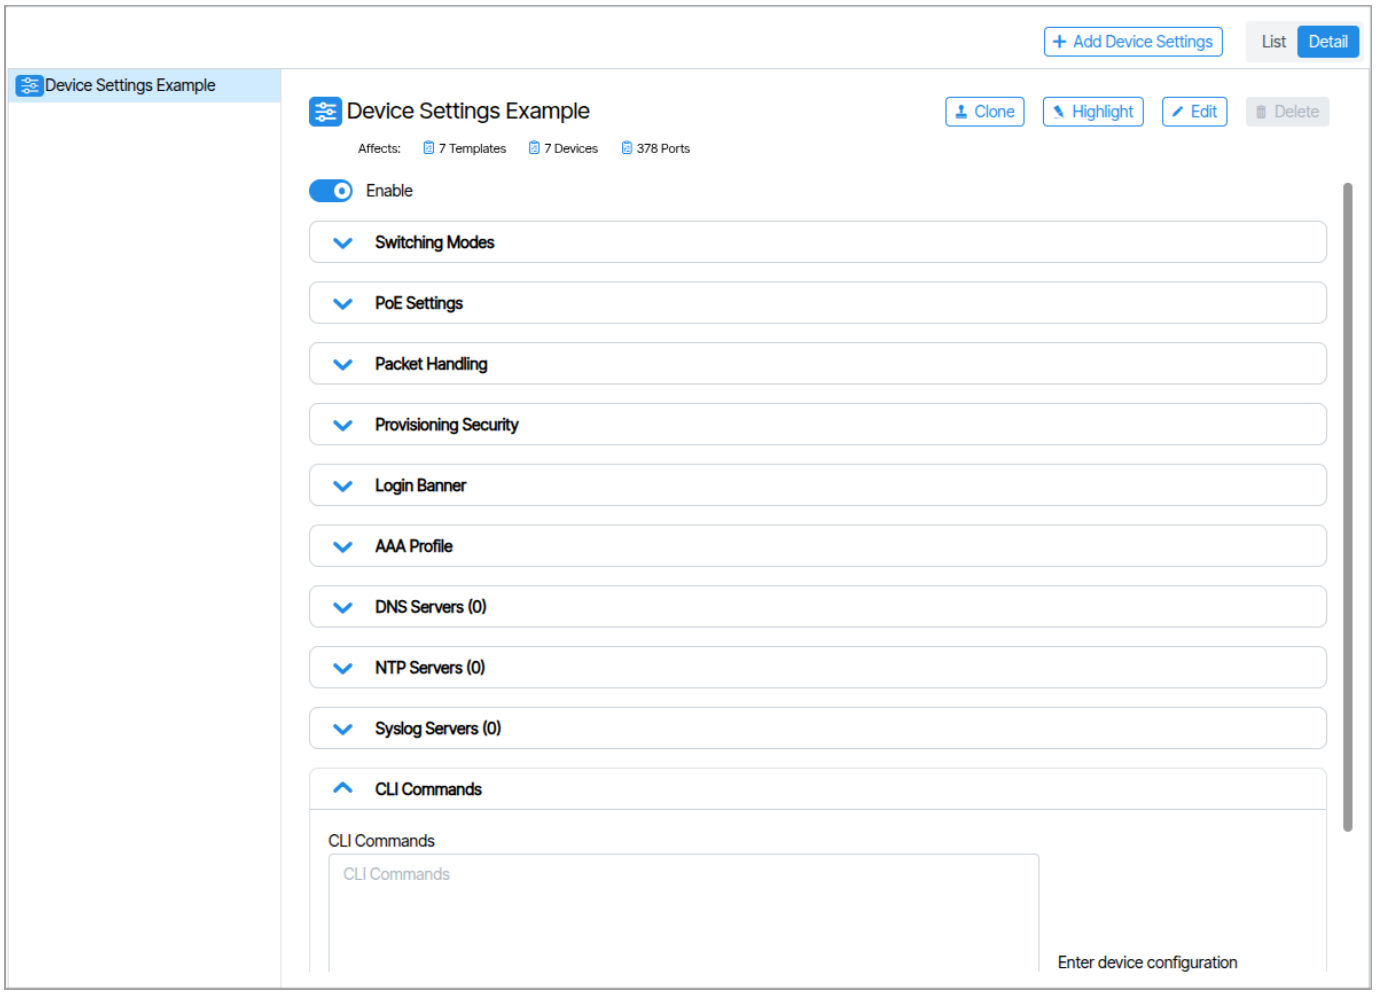
Copyright BE Networks 2026

Device AAA Profile | LLDP | TACACS Profiles

4.5.9 Devices

Device Settings

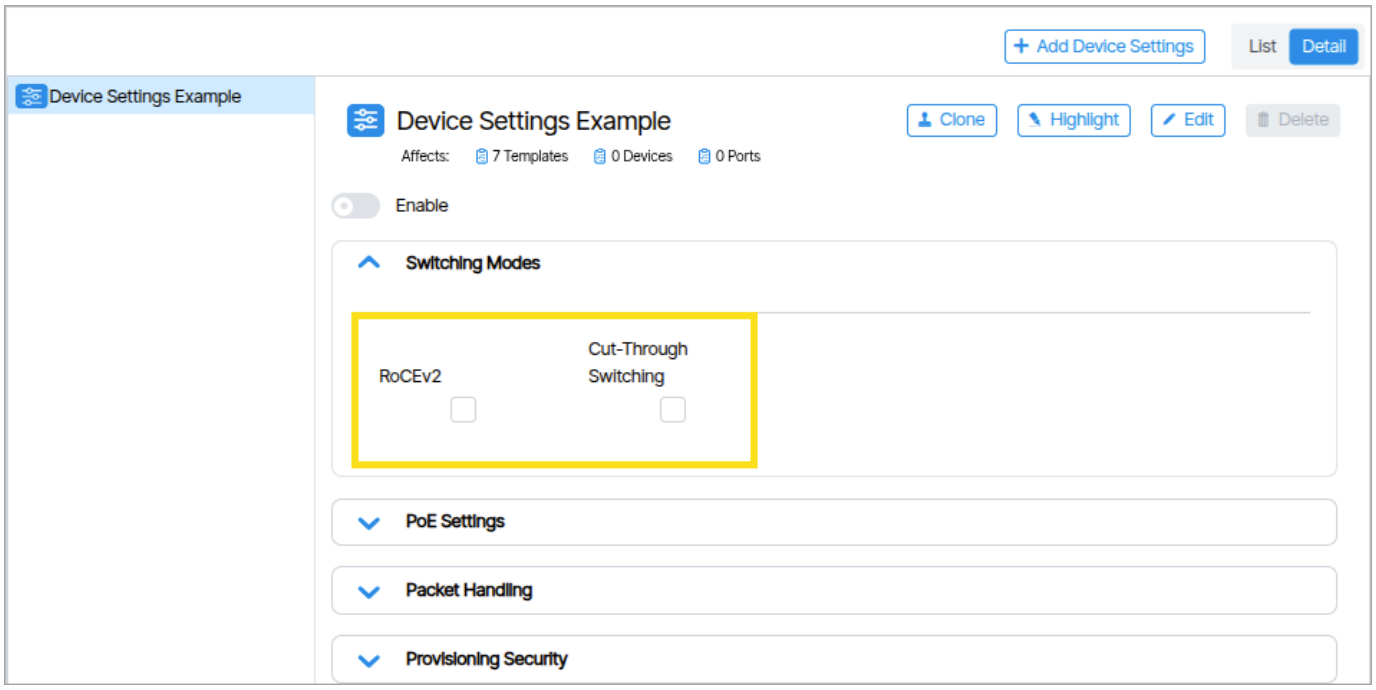
Device Settings are a set of parameters associated with a physical device.



ROCEV2 AND CUT-THROUGH SWITCHING

Device Settings includes settings under the **Network Configuration** field: **RoCEv2** and **Cut-Through Switching**. These options can be enabled or disabled via checkbox selection. Only one of the two selections can be enabled at a time; attempting to enable both items will result in a validation error. Both options can be disabled without triggering a validation error.

Copyright BE Networks 2026



Power over Ethernet Settings

Setting Name	Function
Mode	Automatic or Manual
External Battery Power Available (watts)	When the device is connected to an external battery backup, this is the maximum amount of power that the device should use to power PoE devices. This can be used to force low priority ports off when power is lost to preserve battery for higher priority ports.
External Power Available	Enter the power capability of the power supply powering the device.
Usage Threshold	Percent for comparing the currently consumed power to the allocated power. If the percentage of consumed power over allocated power is over the usage threshold, a fault condition is reported and logged.

Provisioning Security Settings

This functionality encompasses security audit interval configurations that are individually applied to each switch.

Setting Name	Function
Security Audit Interval	Frequency, in minutes, of rereading the switch running configuration and comparing it to expected values. If the value is blank, the audit will use the default switch settings. If the value is 0, the audit will be turned off.
Commit to Flash Interval	Frequency, in minutes, to write the switch configuration to flash. If the value is blank, the commit will use the default switch settings. If the value is 0, the commit will be turned off.

DEVICE AUTHENTICATION PROFILE

The Device Authentication Profile field accepts a **Device AAA Profile**. Assign a Device AAA Profile to configure AAA settings for this device.

TACACS Profiles

A **TACACS Profile** defines the connection parameters for a TACACS+ authentication server, including the server address, shared secret, port, and timeout values. Device AAA Profiles reference it to authenticate and authorize administrative access to network devices.

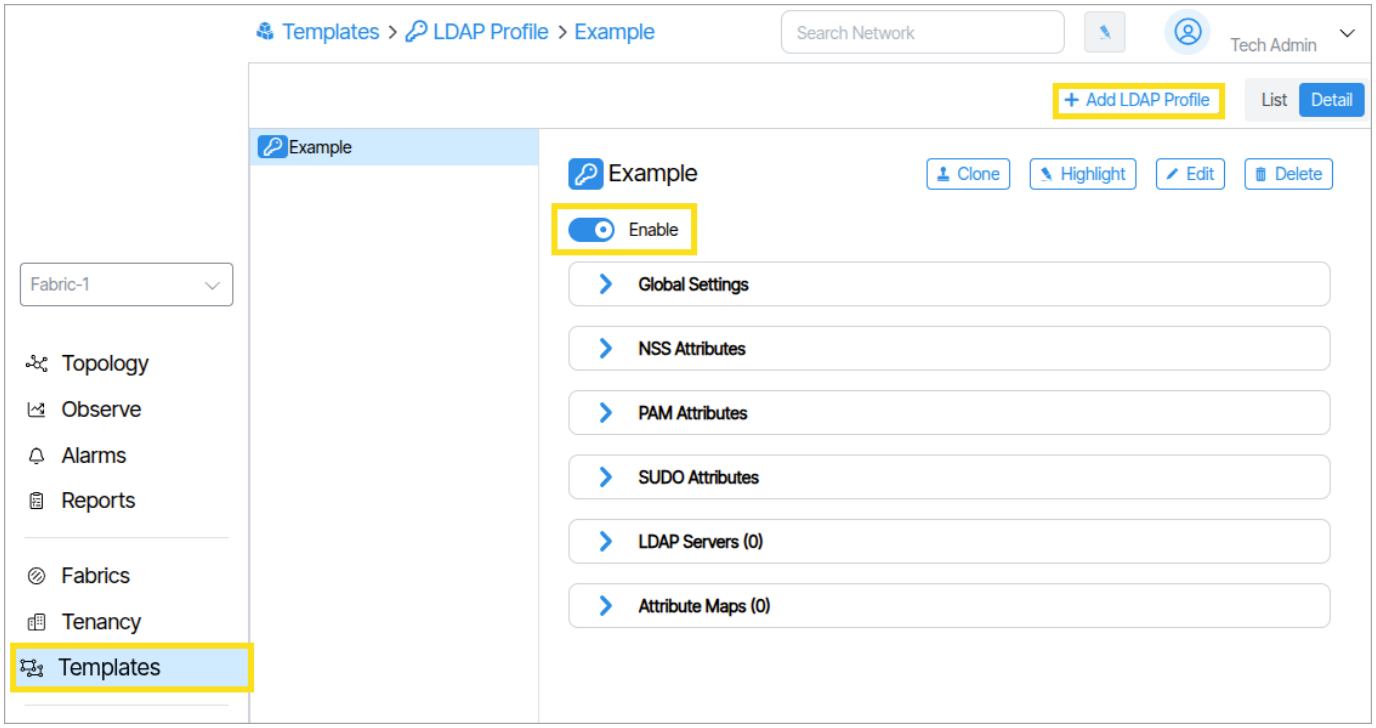
Copyright BE Networks 2026

LDAP Profile

An LDAP Profile is a saved set of LDAP connection and authentication parameters that can be applied to multiple devices without reconfiguring each one individually.

How to Create a Reusable LDAP Profile

1. Go to **Templates → Devices → LDAP Profile**.
2. Click the **Add LDAP Profile** button.
3. Configure the parameters.
4. Enable the profile by clicking the **Enble** switch.



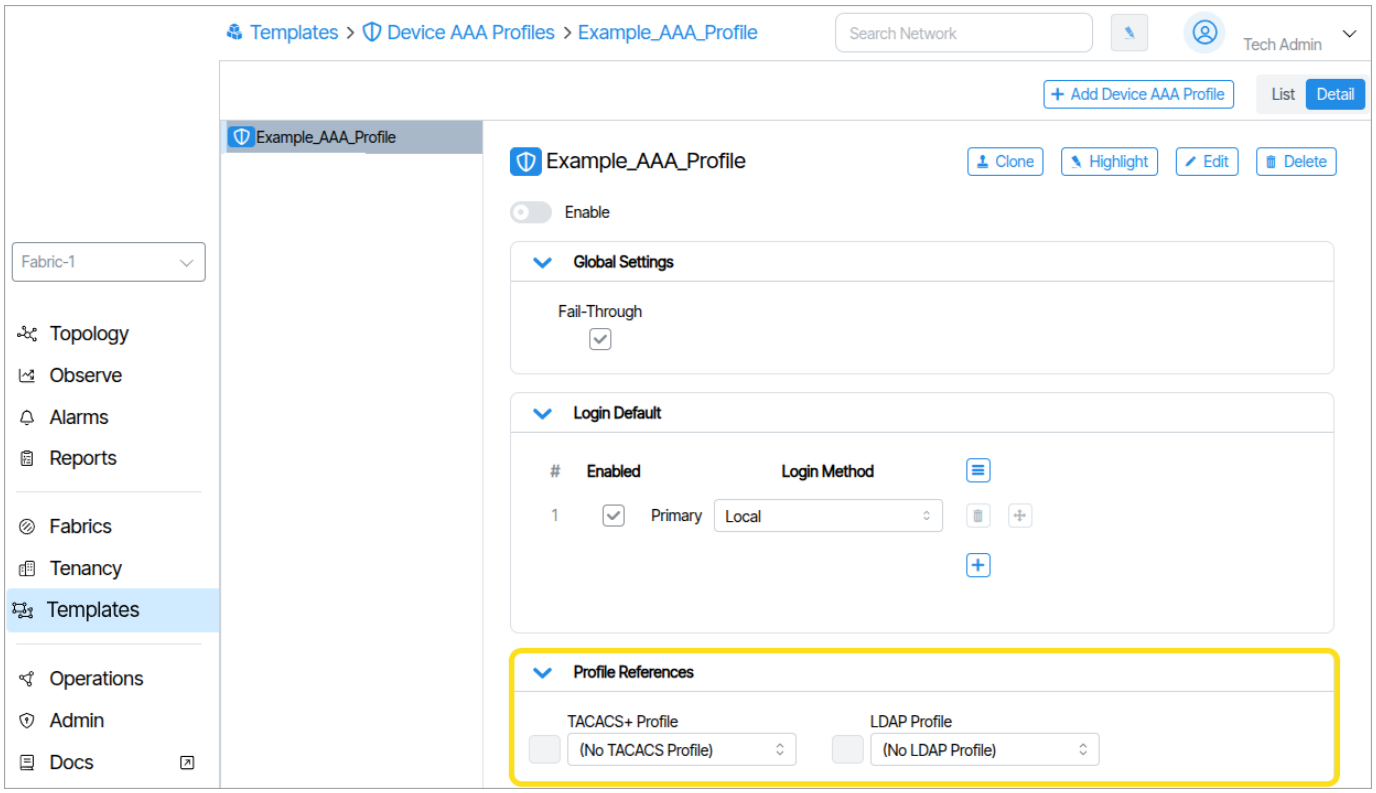
Device AAA Profile

A **Device AAA Profile** defines the authentication, authorization, and accounting policy applied to a network device. It references a TACACS Profile to determine which TACACS+ server handles credential verification, and assigns that policy to one or more devices.

PROFILE REFERENCES

TACACS and LDAP profiles are set at the bottom of the **Device AAA Profile** under the field titled **Profile References**.

Copyright BE Networks 2026



Copyright BE Networks 2026

ERSPAN Monitoring ACL sflow

4.5.10 Diagnostics

Monitoring ACL

Monitoring ACLs select Service VLAN traffic for diagnostic monitoring. A Monitoring ACL references one or more **Services**; the VLANs from those Services are then available to monitoring policies. The Monitoring ACL can be set in **Diagnostic Profiles** and **Traffic Mirrors**.

CREATE A MONITORING ACL

1. Click **Templates**.
2. In the **Diagnostics** section, click **Monitoring ACL**.
3. Click the **Add Monitoring ACL** button to create a new Monitoring ACL.
4. Enter the Monitoring ACL **name**, then create the object.
5. In the **Services** section, add each Service whose VLAN traffic should be selected for monitoring.
6. Toggle the **Enable** switch to enable the object.
7. Click **Save** () in the **Edit Mode** prompt.

A Service may only be added once within the same Monitoring ACL.

MONITORING ACL REPORT

The Monitoring ACL report view shows the configured ACLs and their referenced Services. User-visible columns include **Service Count**, **Service**, and **VLAN**, which can be used to verify which Service VLANs are selected for diagnostic monitoring.

USE A MONITORING ACL WITH A TRAFFIC MIRROR

Traffic Mirrors can reference a **Monitoring ACL** to mirror only the Service VLAN traffic selected by that ACL.

When configuring a Traffic Mirror on a switch endpoint, set the **Monitoring ACL** field to the Monitoring ACL whose Service VLANs should be mirrored. The Traffic Mirror report includes a **Monitoring ACL** column so you can verify which ACL is applied to each mirror entry.

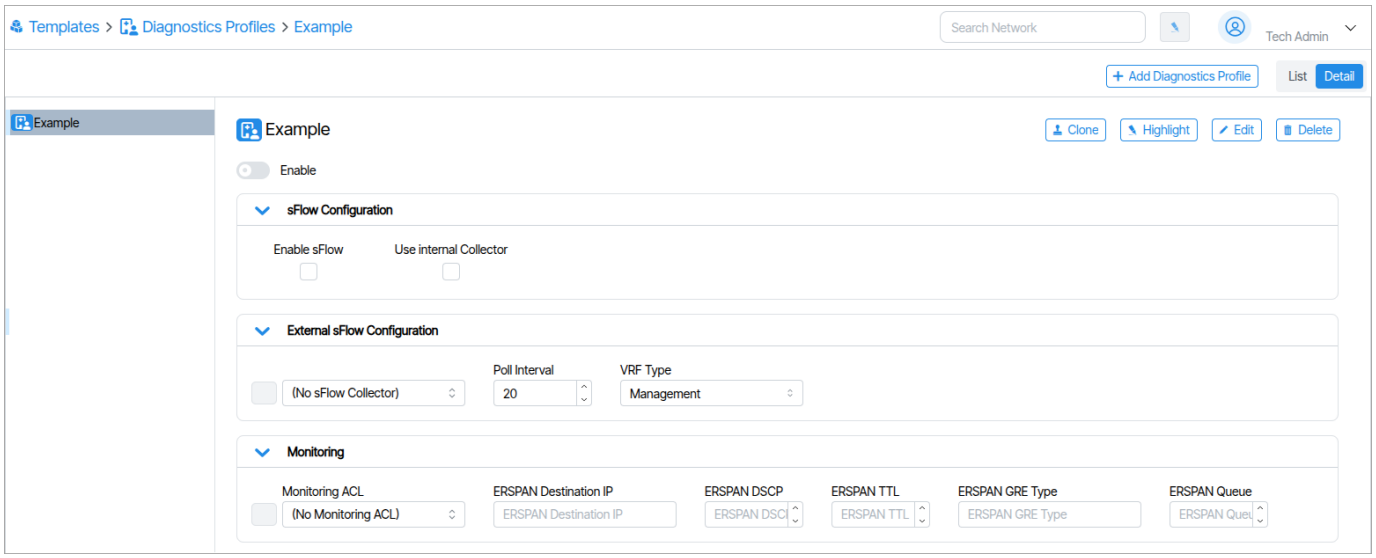
Traffic Mirror source ports also support an **All interfaces** option when a **Monitoring ACL** is selected. When selected, the mirror source is displayed as **All interfaces** in the topology and reports, and CSV import/export uses the value `all` in the **MIRROR SOURCE PORT** column.

When a Traffic Mirror uses a Monitoring ACL, **Outbound Traffic** is not supported. The mirror must use **Inbound Traffic**, and only one enabled Traffic Mirror row on the same switch can use **All interfaces** as the source.

Diagnostics Profiles

Diagnostics Profiles are device level templates that enable and configure sFlow and ERSPAN monitoring on network switches, including settings for collectors, polling intervals, VRF types, Monitoring ACLs, and ERSPAN destinations.

Copyright BE Networks 2026



Diagnostics Port Profiles

Diagnostics Port Profiles are port level templates that control which individual switch ports send sFlow traffic data to the configured collectors.

sFlow

sFlow is a configurable feature in Dell Enterprise SONiC that enables administrators to monitor network traffic using flow based sampling.

> **Feature**
>
> This feature is only available for systems devices running Dell Enterprise SONiC OS.

Copyright BE Networks 2026

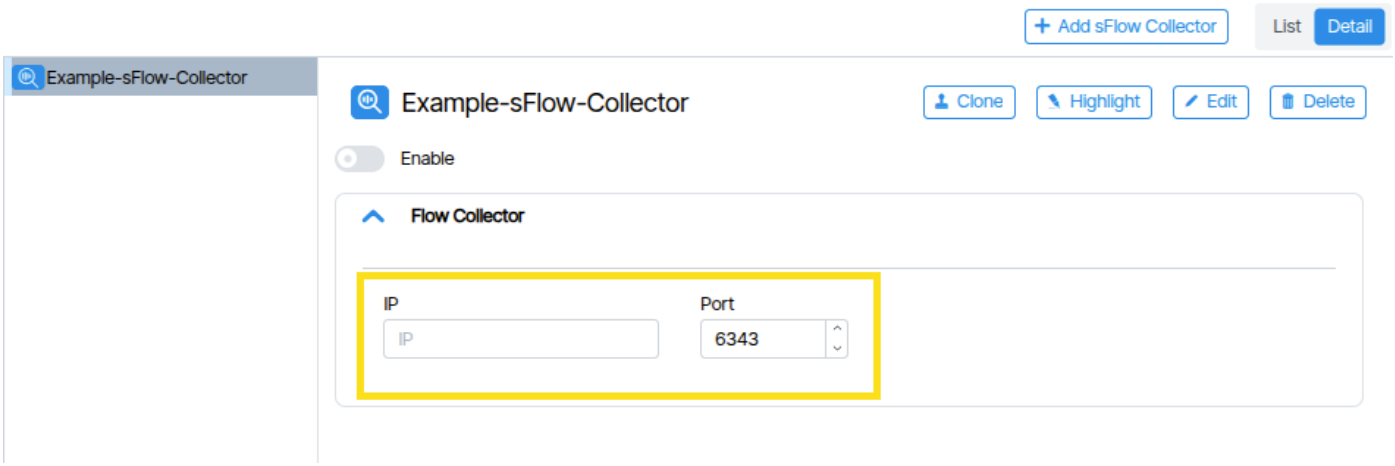
ENABLE SFLOW EXPORTS

Step 1: Configure sFlow Collector(s)

1. Double-click **Templates**.
2. Click **sFlow Collectors**.
3. Click the **Add sFlow Collector** button to create a new sFlow Collector.
4. Enter the collector **name**, then click **Create SFlow Collector**. ()
5. Set:

a. **Flow Collector IP Address**

b. **Port Number** (Default: UDP 6343)



6. Toggle the **Enable** switch to enable the object.
7. Click **Save** () in the **Edit Mode** prompt.

Step 2: Assign Collector to a Diagnostic Profile

1. Double-click **Templates** and click **Diagnostic Profiles**.
2. Click the **Add Diagnostics Profile** button to create a new Diagnostic Profile.
3. In the pop-up window type a name and click **Add Diagnostics Profile**.
4. In the parameters window, configure the following parameters:

a. **Enable**: Toggles the Diagnostic Profile on/off (set it to *enabled*).

b. **Enable sFlow**: Enables sFlow for this profile (set it to *enabled*).

c. **Collector**: Select the desired sFlow Collector (set it to the Collector you created in the previous steps)

d. **Poll Interval**: Interval (in seconds) for traffic sampling or counters. Range: 5–300 (default: 20).

e. **VRF Type**: Choose **Management** or **Underlay**, depending on your network architecture.

Step 3: Assign the Diagnostic Profile to Device(s)

1. In **Topology** from **Network View**, select the switch you want to monitor.
2. Double-click the switch to zoom in.
3. At the top of the device double-click the **Diagnostic Profile** window, assign the appropriate Diagnostic Profile (the one linked to your sFlow Collector).



Copyright BE Networks 2026

Step 4: Configure Ports to Send Data to Collector(s)

1. Double-click **Templates** and click **Diagnostic Port Profiles**
2. Enable the **Diagnostic Port Profile** and Enable **sFlow** .
3. Navigate to the switch port(s) connected to your monitored device or collection path and double-click (or zoom in) on it.
4. Double-click the port's endpoint to open configuration options.
5. In the **Diagnostic Port Profile** menu assign the appropriate **Diagnostic Profile** to the port. 

PB Egress Profiles

PB Egress Profiles are groups of filters applied to Packet Broker egress connections.

Copyright BE Networks 2026

4.6 Observability

Monitoring Observe

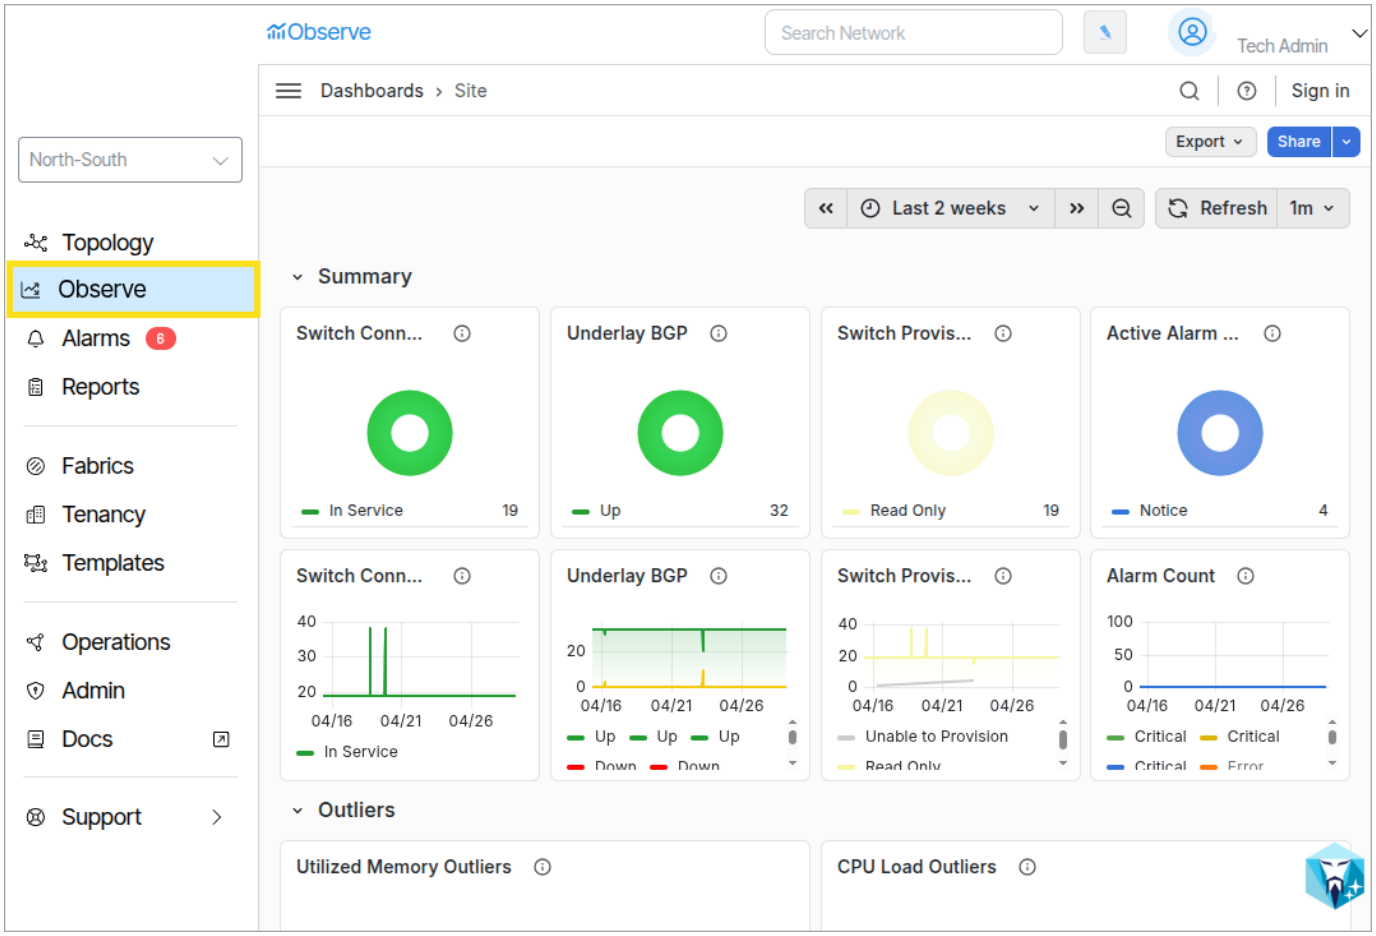
4.6.1 Observe

Installation

For installation instructions go here: Monitoring Installation Instructions For instructions on how to update to the latest Monitoring software go here: Upgrade to the Latest Monitoring Software.Monitoring Observability Functionality is all accessible in the Observe menu option.

Observe is a collection of networking monitoring tools used to track system-wide information, such as service port activity, virtual machine performance statistics, provisioning information, and more.

To access Observe, click the **Observe** option on the Main Navigation Bar.



Copyright BE Networks 2026

OBSERVE DASHBOARD TYPES

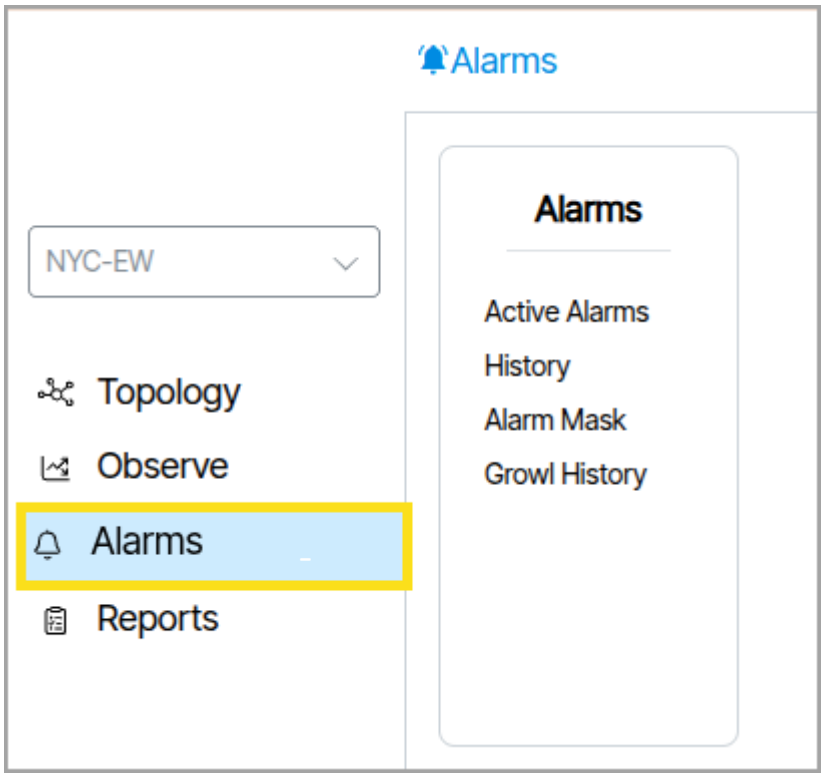
Dashboard	Function
Fabric	The Fabric dashboard presents high-level statistical information for all devices in the network, such as the percentage of in-service devices and general connectivity data.
LAG	The LAG dashboard presents data specific to each LAG. You can select the desired LAG using the LAG dropdown menu.
Pod	The Pod dashboard presents data specific to each pod. You can select the desired pod using the Pod dropdown menu.
Tenant	The Tenant dashboard presents data specific to each tenant. You can select the desired tenant by using the Tenant drop down menu.
Switch	The Switch dashboard presents data specific to each individual switch. You can select the desired switch by using the Switch and Port Type dropdown menus.
Port	The Port dashboard presents data specific to each individual port. You can select the desired port by using the Port drop down menu.
System	The System dashboard presents statistical and graphical data for the high-level system components that make up the system. These components include the VnetC, SDLC, and Monitoring virtual machines, as well as the ACS and Backend Service(s), which are containers within the SDLC.
Executive	The Executive dashboard provides system overview data.

Copyright BE Networks 2026

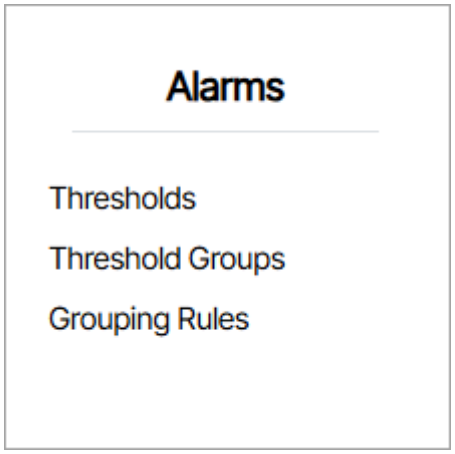
Alarms Monitoring

4.6.2 Alarms

Alarms notify the administrator about problems within the system. The administrator accesses alarm information from the Alarms navigation panel.



The following settings—Thresholds, Threshold Groups, and Grouping Rules—are available in **Templates → Alarms**.



The **Thresholds** feature consists of three sections that work together: **Thresholds**, **Threshold Groups**, and **Grouping Rules**. Together, these create automated notifications that alert administrators when specific metrics exceed, fall below, or meet certain conditions in your network infrastructure.

Thresholds

Thresholds allows administrators to create multiple threshold rules, each monitoring a specific metric with configurable conditions (greater than, less than, equals, etc.) and severity levels.

Copyright BE Networks 2026

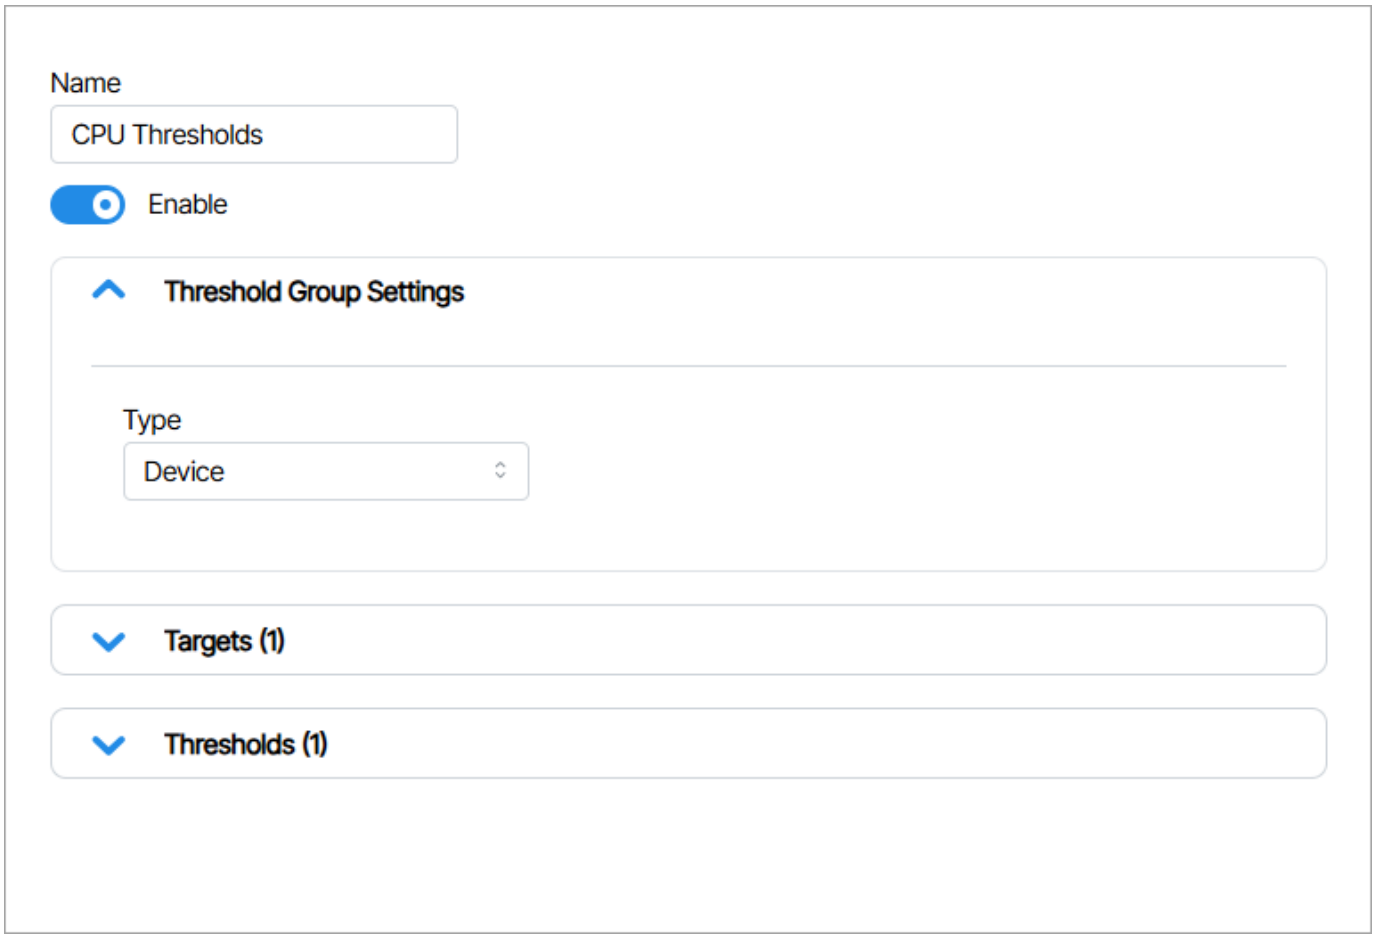
Grouping Rules

Grouping Rules are collections of logical rules that determine which ports and hardware devices to select.

Threshold Groups

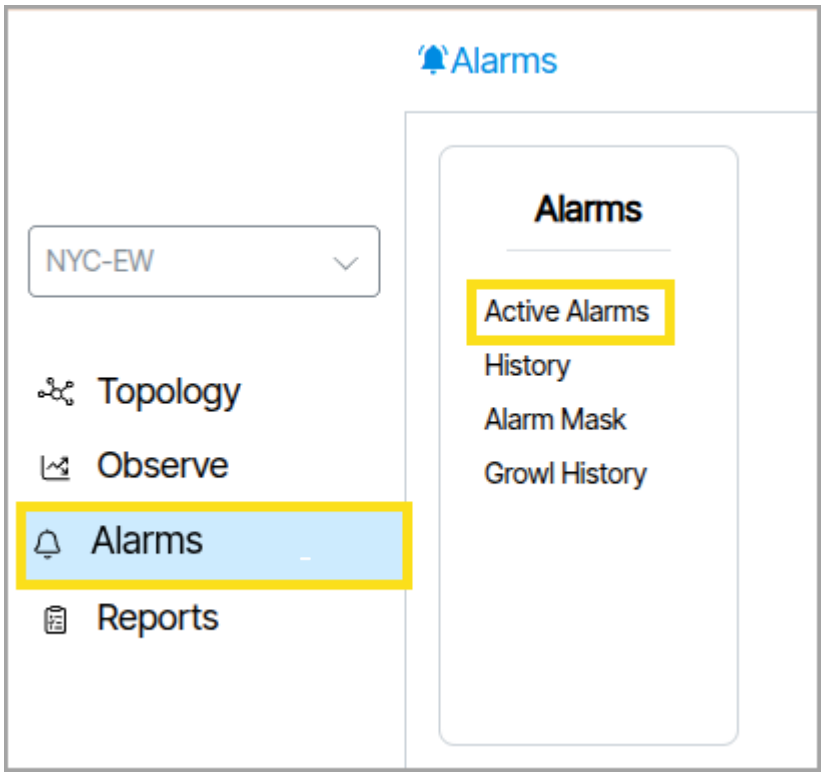
Threshold Groups is where metric watching is applied to devices and ports. This section is composed of two lists: Targets and Thresholds.

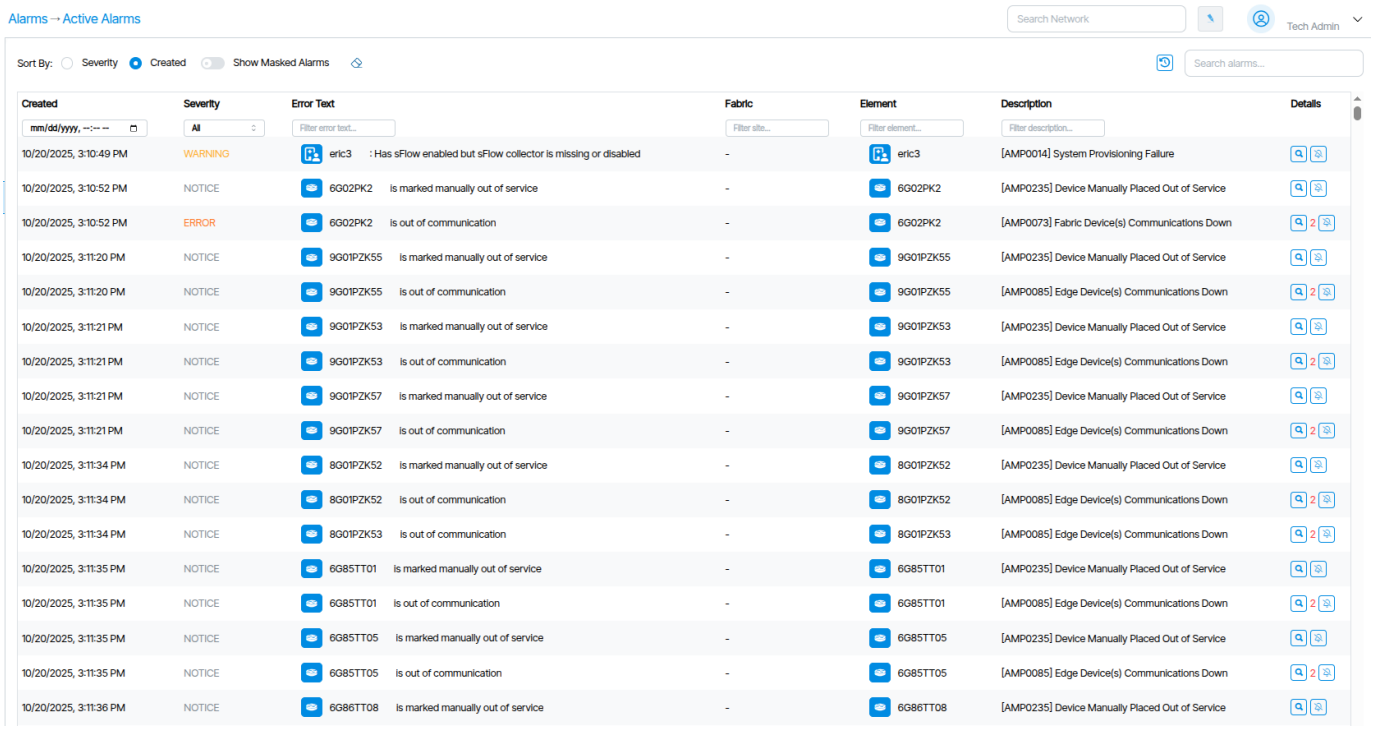
- **Targets** is where you assign the rules created in the Grouping Rules section.
- **Thresholds** is where you assign the rules created in the Thresholds section.



Active Alarms convey immediate problems occurring in the network right now. The Active Alarms window displays these alarms. A red background highlights the number of active alarms next to the Alarms Main Navigation tab. 

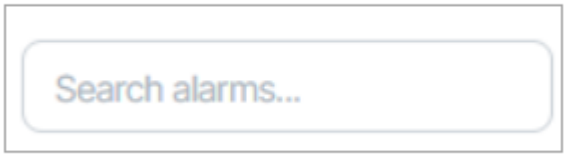
Copyright BE Networks 2026



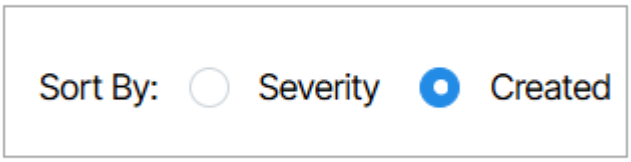


Active Alarms are searchable by name and object type using the search field.

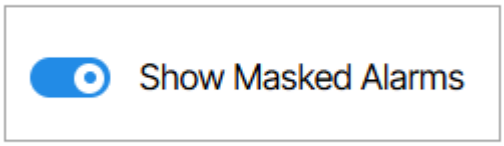
Copyright BE Networks 2026



The **Sort By** feature is used to sort alarms. Alarms are sorted by **Severity** or by their creation time by enabling the **Created** option.

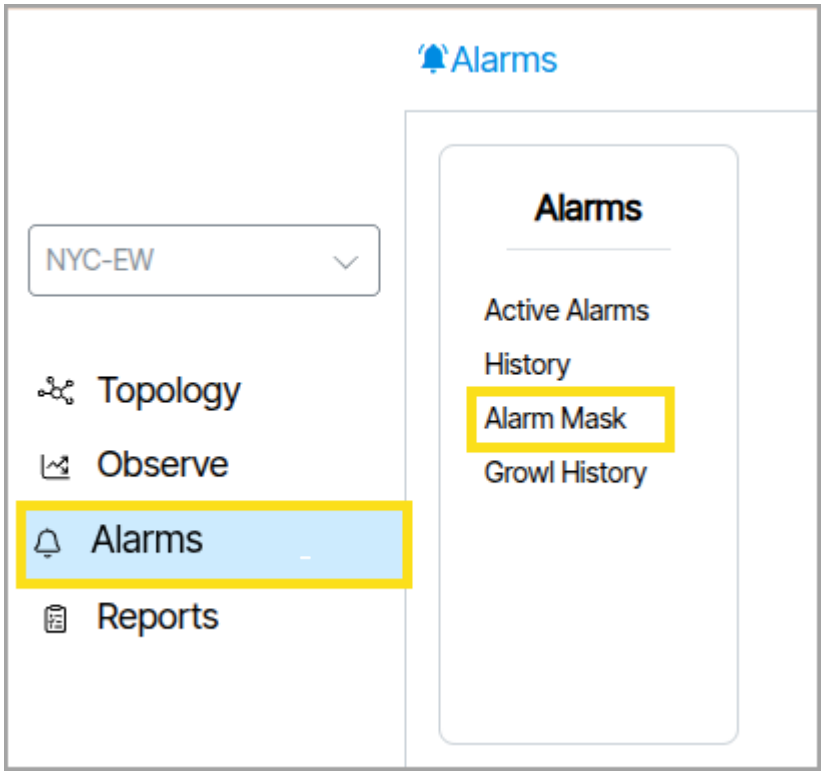


The **Show Masked Alarms** option shows **Masked Alarms** from within the Active Alarms window.

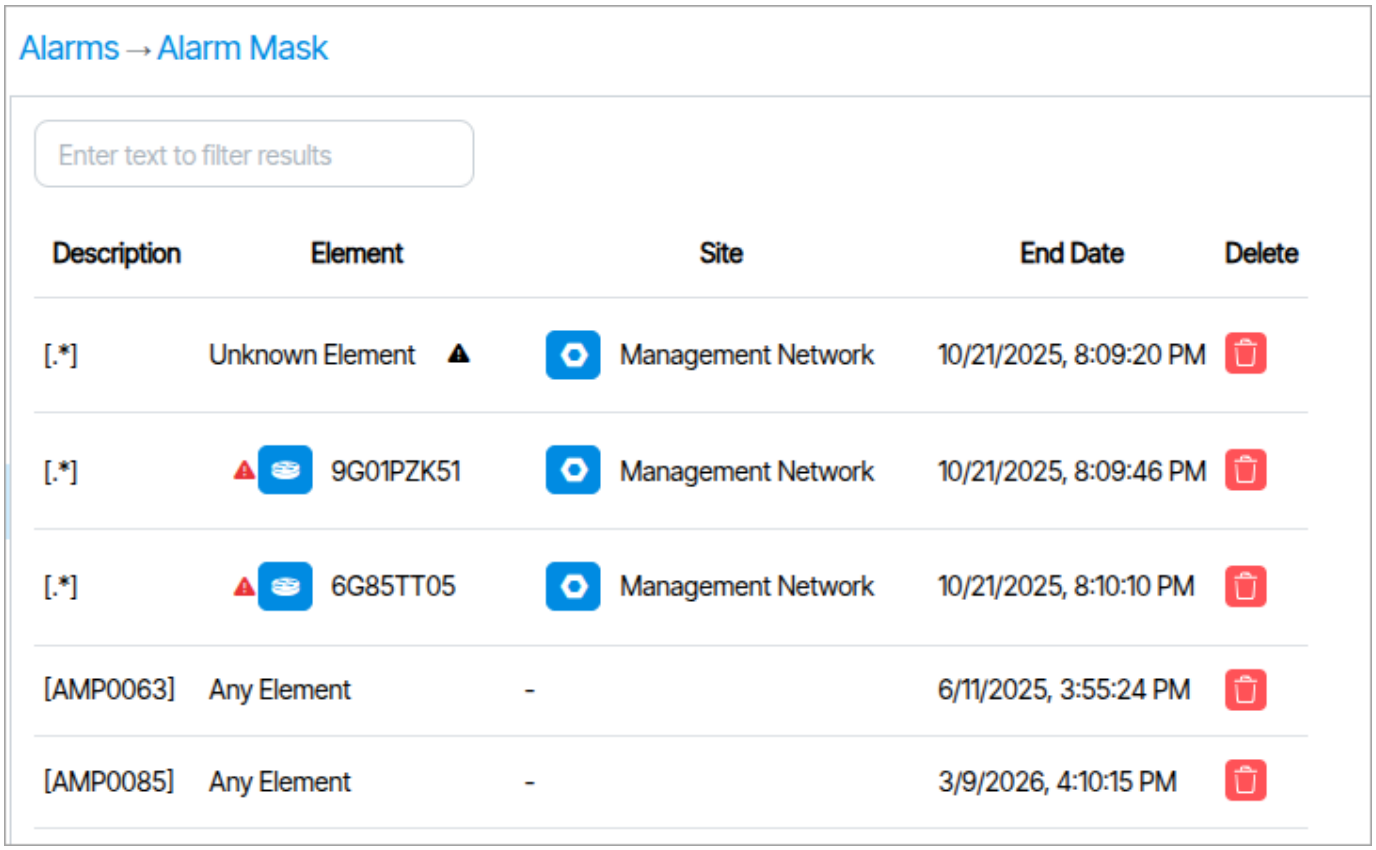


4.6.3 Alarm Mask

Alarm Mask lists Masked Alarms.

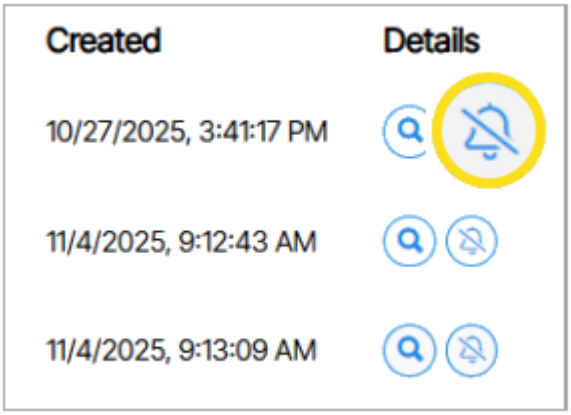


Copyright BE Networks 2026



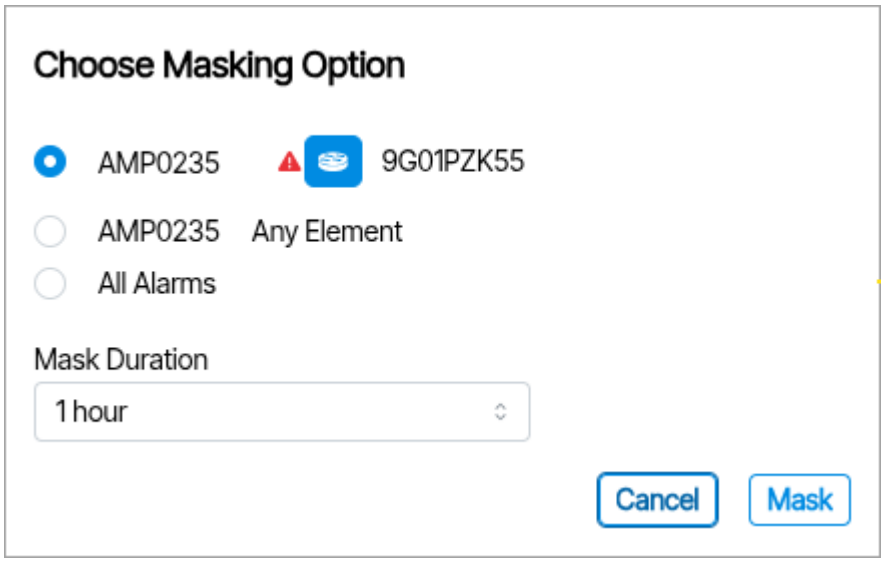
Masked Alarms

Masked Alarms hide from the Active Alarms list. To mask an alarm, click its Mask Alarm button under the **Details** column in the Active Alarms window.



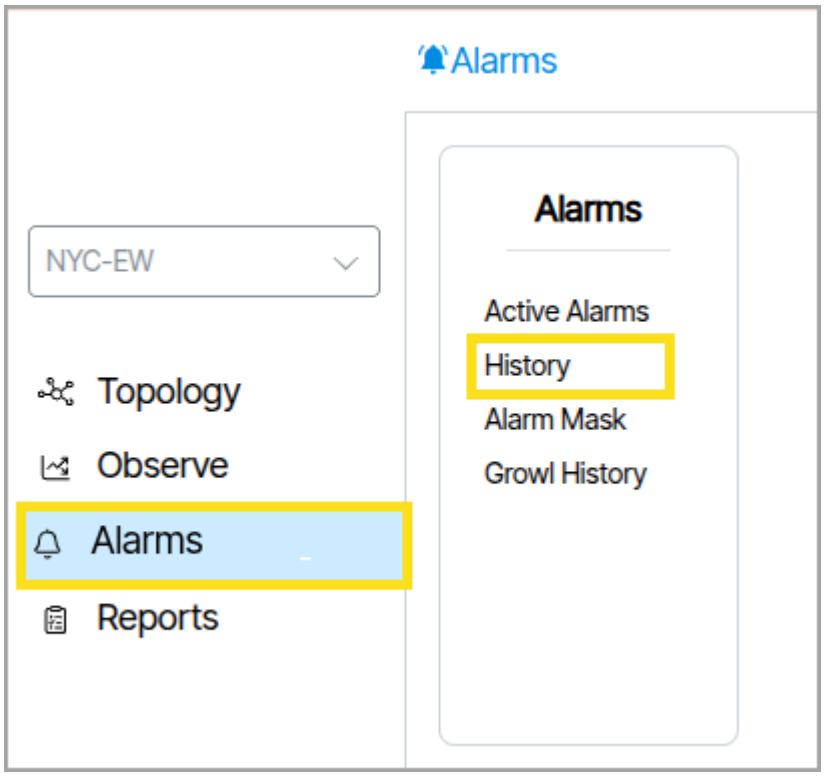
When you mask an alarm, a pop-up box appears with additional options, including **Mask Duration**, which lets you set a time duration for the alarm mask.

Copyright BE Networks 2026



4.6.4 Alarm History

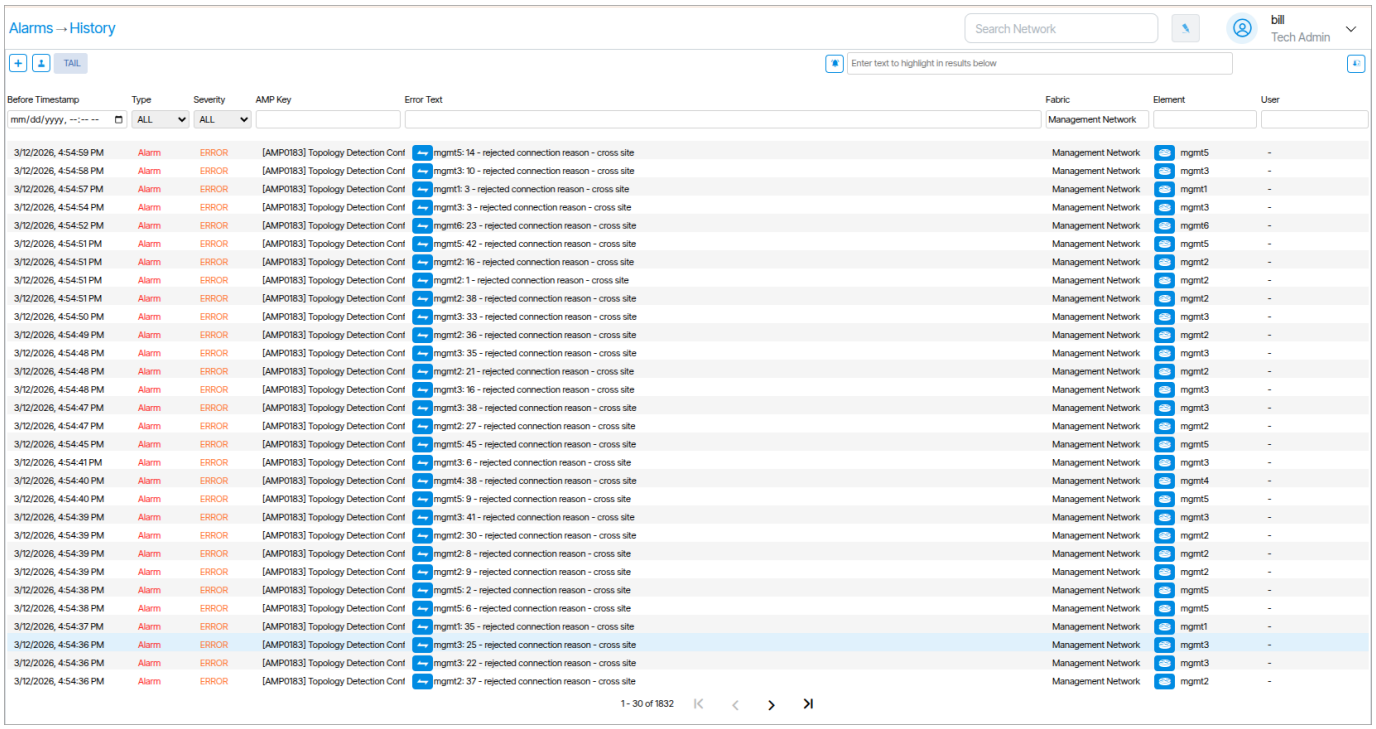
While **Active Alarms** notify administrators of immediate problems, **History** records previously resolved issues and tracks less pressing notifications such as user logins.



Columns

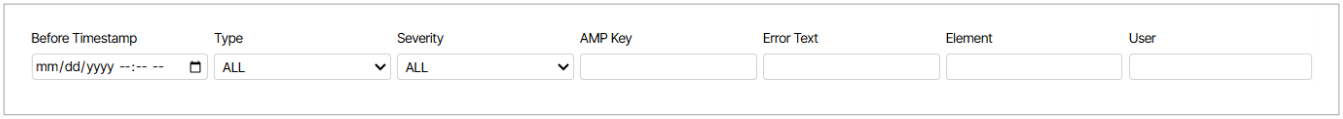
Each column showcases data pertinent to individual alarm events, with a corresponding description positioned at the top.

Copyright BE Networks 2026



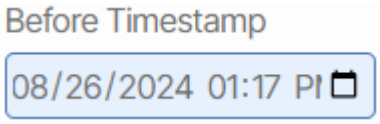
Filtering

Each column features a form field positioned above it. These form fields enable the filtering of alarm data based on the content of the column. To use the filters, you choose a field(s), input a query, and press **Enter** on your keyboard to apply the filter. In addition to conventional search text, filters also accept Regex and glob patterns.



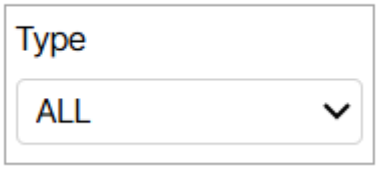
Before Timestamp

Before Timestamp excludes events that happened after a set date. After the fields are set, only events that happened before the selected time value are displayed.



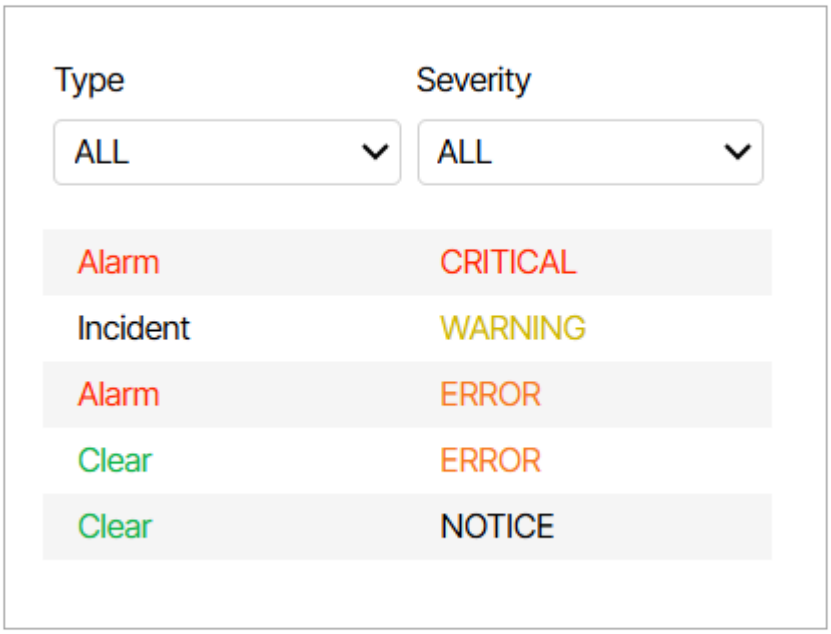
Type

Type filters alarms by their type, which can be **Alarm**, **Event** or **Clear**. Setting the value to **All** shows all items and invalidates the filter.



Messages under Type and Severity are color-coded by **Type**.

Copyright BE Networks 2026



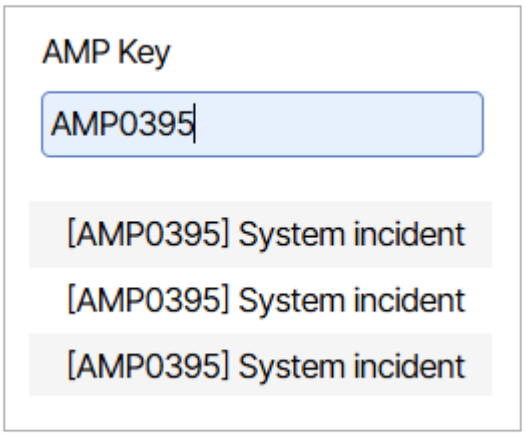
Severity

Severity filters alarms by their severity, such as **Critical**, **Error**, **Warning**, or **Notice**. Setting the value to **All** displays all items and disables the filter.



AMP Key

The **AMP Key** filters alarms by key. The key is placed in brackets at the beginning of the alarm description.



Copyright BE Networks 2026

Error Text

The **Error Text** field filters alarms based on the content of **Error Text** column.



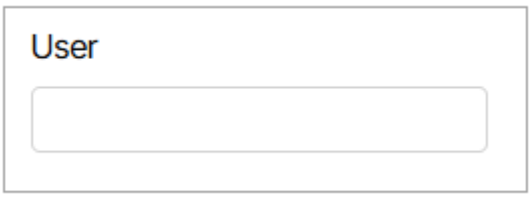
Element

Element filters alarms by device name.



User

The **User** column indicates who triggered an alarm. This field filters alarms by the **User** responsible for them.



Site

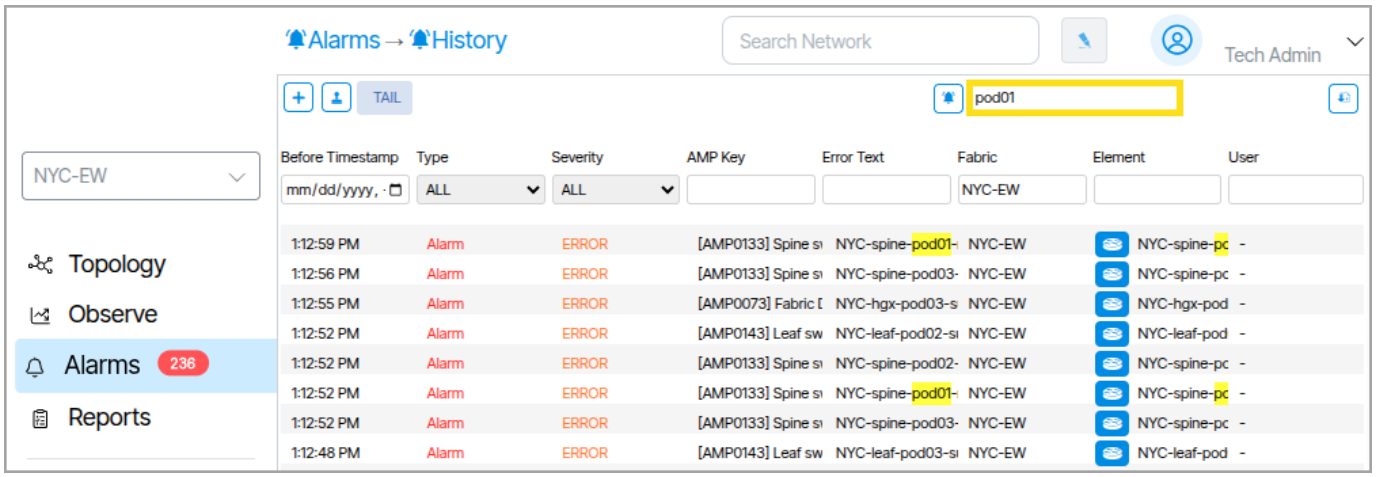
> **Multisite**
>
> This column is only available on *multisite* systems.

The **Site** column indicates the site of the device.

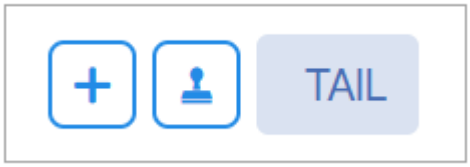
Search

The search feature at the top of the page allows you to search for any text displayed in the current list of alarms. Any matching text is highlighted in yellow.

Copyright BE Networks 2026



Custom Search Filters

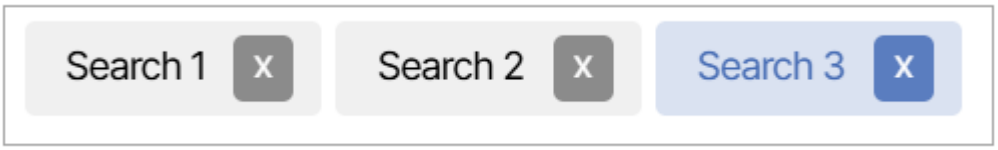


The **TAIL** tab lets you set the initial state of all filter field values. This is useful because it enables you to create a collection of default filter settings before creating custom search filters.

Add a Custom Search Filter

The + button allows you to create a custom search filter. This is a user-defined collection of all filter settings. Multiple custom search filters can be created and toggled between when querying Alarm data.

When custom search filters are created, they are represented as numbered tabs with the word *Search* displayed on them. Clicking these tabs anywhere other than on the displayed "x" selects them and populates the fields with the assigned values. Clicking the "x" deletes the search template and removes it from the list. Clicking **Tail** after creating or selecting a custom search will revert the filter values to the state saved via the **Tail** tab.



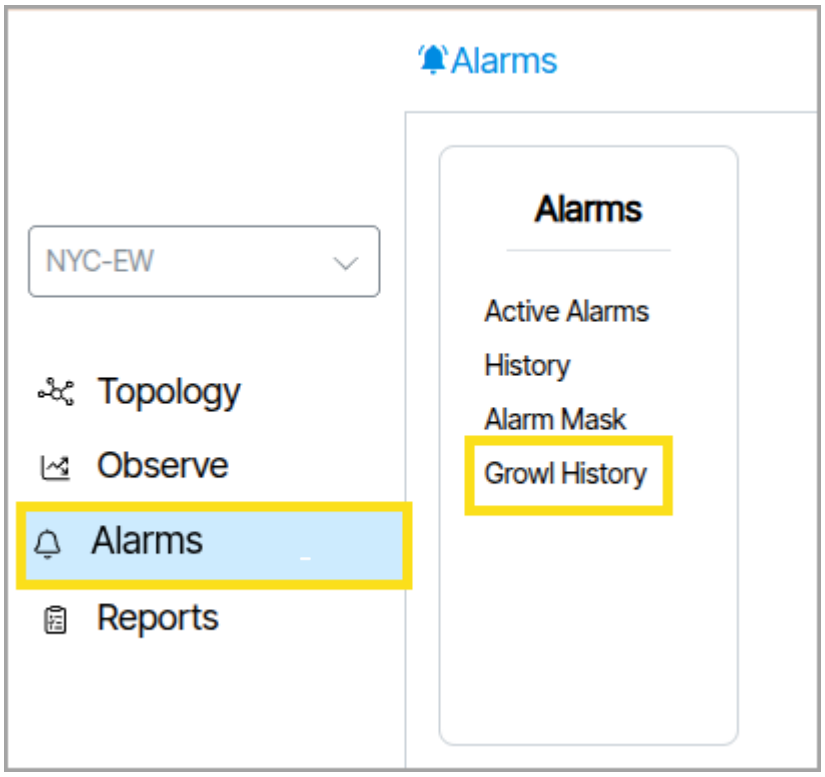
Copy a Custom Search Filter

The button allows you to clone all filters from the currently active view window. The clone is presented as a newly created active tab.

4.6.5 Growl History

Growl History lists all system Growls. Growls are errors that occur between the client browser and the Verity orchestration platform. Hovering the mouse over a growl displays more detail about the error.

Copyright BE Networks 2026



Copyright BE Networks 2026

Monitoring

4.6.6 Alarm List

SNMP MIBs available here.

Dict key	Severity	Type	Description
AMP0013	ERROR	alarm	System Provisioning Failure

Element type	Error text	Corrective action
Gateway	{{ element_id }} has the same {{ fields }} as {{ gateway_id }} while provisioned on the same switch	Correct the provisioning conflict

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0014	WARNING	alarm	System Provisioning Failure

Copyright BE Networks 2026

Element type	Error text	Corrective action
Gateway	{{ element_id }} is provisioned multiple times on {{ switch_id }}	Correct the provisioning conflict
Gateway Profile	{{ element_id }}'s Virtual Port VLAN conflicts with {{ conflict_id }}'s VLAN	Correct the provisioning conflict
Gateway Profile	{{ element_id }} has more than one BGP Gateway with an empty Egress VLAN	Only one BGP session can configured on a L3 configured port
LAG	{{ element_port }} - {{ vlan }} has conflicting external VLAN	Correct the invalid Ethernet Port Profile
LAG	{{ element_port }} - is Peer-LAG/IDS and provisioned across multiple sites	Switch pairs must exist on the same site. Either correct the site assignment or correctly identify the switches intended to be in a pair.
Endpoint Bundle	{{ element_id }} has one or more eth port profiles with conflicting RADIUS names.	Update the Profile to remove the conflict
Device Controller	{{ element_id }} LLDP chassis identifier does not match managed device	Update the chassis or serial number identified correct in the Device Controller
Device Controller	{{ element_id }} sharing Management VLAN with ONTs, cannot be Managed on native VLAN	The service in the Active SFP Device Controller must be changed to use a different management VLAN than the one used by the ONTs connected on the OLT.
Site	{{ element_id }} has multiple services using VLAN {{ vlan }} with conflicting settings	It is not recommended to use the same VLAN ID on the same site. Change the VLAN IDs to be unique for the site.
Site	{{ element_id }} missing management service	Ensure that site settings are referencing a valid management service
Site	{{ element_id }}: {{ service_id }} is Externally Routed but has no Service Ports	Services may switch only locally inside the LAN, or can connect to the uplinks of the LAN to egress the network. The option in the service indicates whether to expect this service as local only or leaving the fabric. In this case, the service is marked to egress the fabric, but the service is not provisioned on any uplink Service Port Profiles. Correct the provisioning in either case.
Switch	{{ element_id }} connected to SD-WAN is not provisioned as TOR	Remove {{ element_id }}'s connection to SD-WAN or provision as TOR
Site	Base Management VLAN not defined	Define in Admin - Settings tab
Device Controller	{{ element_id }} is provisioned on {{ prov_device_port }} but it's managed Device is appearing on {{ real_device_port }}	Check the network fabric connections to make sure that the provisioning of the switch controller matches the desired network physical connectivity. Correction required either in provisioning or cabling.
Device Controller	{{ element_id }} is provisioned as Static Port, but the location is incomplete	Finish provisioning the location
Device Controller	Discovery non-operational because no Site is defined	Contact support and provide system snapshot
Device Controller	LLDP Search String does not match the device currently connected to {{ device_port }}	Check the network fabric connections to make sure that the provisioning of the switch controller matches the desired network physical connectivity. Correction required either in provisioning or cabling.
Endpoint Bundle		Fix protocol mismatch

Copyright BE Networks 2026

Element type	Error text	Corrective action
	{{ element_id }} has mismatched voice protocol with Voice Device Profile {{ voice_device_profile_id }}	
Endpoint Bundle	{{ element_id }} has mismatched voice protocol with Voice Port Profile {{ voice_port_profile_id }} on POTS ${row}	Fix protocol mismatch
Endpoint Bundle	{{ element_id }} - the RG feature will not be operational	Fix provisioning
Endpoint Bundle	{{ element_id }} - IPTV cannot be processed two ways. IPTV application (IGMP Proxy) and IPTV Filtering (IGMP Snooping) both provisioned on single ONT, IGMP Proxy will take precedence	Remove the IPTV application or the IPTV Filtering
Endpoint Bundle	{{ element_id }} - Port advertised by LLDPmed is not provisioned in the Eth-Port Settings and is not present in the Eth Port Profile	Provision the advertised port or correct the assigned Ethernet Port profile.
ONT	{{ element_id }} - LLDP MED service {{ service_id }} is being used on a translated 802.1x port	Remove the service. These features are mutually exclusive.
Eth-Port Profile	{{ element_id }} has VLAN collision on VLAN {{ vlan }} and {{ acl_id }}. Rows cannot have the same MAC filter and exit VLAN	Remove one of the conflicting VLANs
Eth-Port Profile,ONT,Switch	{{ element_id }} has a VLAN collision on VLAN {{ vlan }}.	Remove one of the conflicting VLANs
Switch,Eth-Port Profile	{{ element_id }} has multiple services with the same VLAN {{ vlan }}.	Remove one of the conflicting services
Eth-Port Profile,Service Port Profile	{{ element_id }} - TLS trunk with a non-TLS service	Remove the non-TLS service or change this to not be a TLS trunk
Interface	MAC Security provisioned but will not be applied to {{ element_port }}	Fix MAC collision
LAG	{{ element_id }} is a IDL/IDS LAG whose provisioned ports do not match it's physical connections.	Fix the provisioning/physical mismatch
LAG	{{ element_id }} is IDL/IDS and provisioned across multiple Sites.	Multichassis Switch pairs cannot cross site boundaries. Limit provisioning to a single site
LAG	{{ element_id }} is missing physical connections on provisioned fabric LAG ports	Provisioning indicates an expected connection from downstream device. Connect the new device or unprovision the port
LAG	{{ element_id }} has mismatched interfaces, Upstream : {{ up_size }}, Downstream : {{ dn_size }}	Correct the Port Settings of either the upstream or downstream interfaces
Switch	{{ element_id }} - OLAG detected but not supported	Remove the OLAG
Switch	{{ element_id }} has more than the supported number of LAGs. Supported:	Remove extra LAGs

Copyright BE Networks 2026

Element type	Error text {{ supported_count }}, Current: {{ current_count }}	Corrective action
Switch	{{ element_id }} Mistmached provisioning on TOR uplinks	Make sure the provisioning is the same on all uplinks
Interface	{{ element_port }} PoE allocated is greater than PoE available on {{ element_id }}	Reduce allocated PoE
Interface	{{ element_id }} - {{ prov_error }}	Fix the provisioning or connections
Static IP	{{ element_id }} - {{ service_id }} not provisioned on targeted {{ located_by }}	Provision the service on the correct location
Service	{{ element_id }} is provisioned on too many leaf swithes	The service is provisioned on more leaf swithes than its Tenant's DHCP relay subnet supports. Increase the subnet size of the DHCP relay subnet so that more addresses can be allocated for the Tenant.
Service	{{ element_id }} has invalid DHCP server provisioned.	Service has DHCP server IP without a DHCP Relay source subnet defined for the Tenant.
Switch Pair	{{ element_pair }} no interconnect provisioned	Provision an interconnect
Interface	{{ element_port }} has {{ eth_id }} provisioned but, {{ radius_id }} has no enabled rows	Add an enabled Radius entry
Endpoint Bundle	{{ element_id }} - mixed (trusted and untrusted) port levels are not supported on ONTs. All Eth Port Profiles must be of one kind. By default, all ports will now be untrusted.	Change to use only trusted or untrusted port levels
Endpoint Bundle	{{ element_id }}, Authenticated Eth-Port provisioned but RADIUS is disabled or missing a service	Enable RADIUS and add a service
Switch	ACS {{ element_id }} doesn't have IP Source type set	Complete provisioning of the ACS device.
Switch	Backend System(s) object was not found, skip provisioning	Ensure installation process successfully compeleted and that an SDLC with ACS and Backend System(s) has been created.
Switch	Managed device {{ element_id }} does not have an associated Device Controller	Correct provisioning of the Device Controller if this physical device is expected to be on the network.
Switch	Maximum number of Device Controllers reached for device {{ element_id }}	More SDLC VMs should be added to the system to handle the excess capacity.
Switch	ACS object was not found	Ensure installation process successfully compeleted and that an SDLC with ACS and Backend System(s) has been created.
Switch	Site {{ site_id }} IDL {{ idl_id }} doesn't have both the switchpoints listed	Ensure that the switch pairs are properly provisioned with a LAG created and provisioned between them.
ONT	System has no RADIUS server set up but has Authenticated Eth Port profiles provisioned on device {{ element_id }}	Add the provisioning for the RADIUS server in the Admin section or remove the Authenticated Eth Port profile from the bundle.

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0015	NOTICE	alarm	System Provisioning Failure

Element type	Error text	Corrective action
Interface	{{ element_port }} is provisioned and being used as a port mirror destination	Remove from port mirror destination or remove the provisioning
Interface	{{ gateway_profile_id }}'s Source IP/Mask is /30 or greater	Make sure the Source IP/Mask is /29 or less to allow for 3 valid ip addresses
Interface	Source IPs, Neighbor IP address, and Helper Hop IP Address are not unique for {{ gateway_id }} in {{ gateway_profile_id }} and {{ other_gateway_profile_id }}	Make sure all IP's are unique
Interface	{{ gateway_id }}'s Helper Hop IP Address is not within the Source IP/Mask subnet defined in {{ gateway_profile_id }}	Make sure the Helper Hop IP Address is within the Source IP/Mask subnet
Interface	{{ gateway_id }}'s Neighbor IP address is not within the Source IP/Mask subnet defined in {{ gateway_profile_id }}	Make sure the Neighbor IP address is within the Source IP/Mask subnet
Interface	{{ gateway_id }} Source IP/Mask's have different subnets in Switch Pair	Make sure the Source IP/Mask on each profile have the same subnet
Interface	{{ gateway_id }} has the same Peer GW value on boths side of Switch Pair	Change one of them so that they are different
Interface	{{ lagg_id }} has {{ gateway_id }} provisioned on only one side of Switch Pair	Either remove {{ gateway_id }} or add it to the other side

Dict key	Severity	Type	Description
AMP0022	CRITICAL	alarm	Ethernet Port is Down

Element type	Error text	Corrective action
Interface	{{ element_port }} Critical port is link down	Determine cause and correct faulty cables or SFPs. May be transitory alarm as the connected switch reboots. Disable and re-enable the port

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0023	ERROR	alarm	Ethernet Port is Down

Element type	Error text	Corrective action
Interface	{{ element_port }} Storm : Shutdown	Disconnect the device connected to the port and disable and re-enable the port. If problem persists replace the connected device or adjust the storm control settings if that traffic is expected.
Interface		Disconnect the device connected to the port and disable and re-enable the port. If problem persists replace the connected device or adjust the storm control settings if that traffic is expected.
Interface	{{ element_port }} MAC {{ violation }}	Disconnect the device connected to the port and disable and re-enable the port. If problem persists replace the connected device or adjust settings if that MAC is expected.
Interface	{{ element_port }} Unidirectional Connection Detected	Caused by multimode fiber incorrectly connected. Ensure that the fiber connection TX and RX for the same ports are properly connected. Disable and re-enable the port
Interface	{{ element_port }} is disabled by CRC error detection	Caused by excessive connection layer errors. Determine cause and correct faulty cables or SFPs. Also, the CRC error threshold can be adjusted to relax the port removal. Disable and re-enable the port
Interface	{{ element_port }} Fabric port is link down	Determine cause and correct faulty cables or SFPs. May be transitory alarm as the connected switch reboots. Disable and re-enable the port
Interface	{{ element_port }} Monitored port is link down	Determine cause and correct faulty cables or SFPs. May be transitory alarm as the connected switch reboots. Disable and re-enable the port
Interface	{{ element_port }} Loop Detected	Determine cause and correct looped cables. May be directly connected cables or may be caused by unmanaged devices connected off of the Ethernet port. Once the loop is removed, disable and re-enable the port

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0063	ERROR	alarm	Device Provisioning Failure

Element type	Error text	Corrective action
Switch	{{ element_id }} - Device controller failure - {{ status_error }}	Can be the result of communication errors to the device, or a provisioning failure between Verity and the device or and error from the device's operating system. Show the error output dialogue at the top of the switch in the GUI to determine if its a direct error from the OS. If not, try to revert the last config changes of the system to determine if that is the cause. Enable or disable provisioning objects related to last change. If issue persists, capture a system snapshot in the support menu, and a snapshot for the affected switch.
Switch	{{ element_id }}: Device Controller is out of communication, unable to provision	May be a transitory error. If it persists, make sure the connections/routes are intact between the SDLC VM and the Vnetc and that the firewall settings on the vnetc are correct in the Admin Settings page.
ONT,Switch	{{ element_id }} has failed to provision - {{ error }}	Can be the result of communication errors to the device, or a provisioning failure between Verity and the device or and error from the device's operating system. Show the error output dialogue at the top of the switch in the GUI to determine if its a direct error from the OS. If not, try to revert the last config changes of the system to determine if that is the cause. Enable or disable provisioning objects related to last change. If issue persists, capture a system snapshot in the support menu, and a snapshot for the affected switch and send to technical support.
ONT	{{ element_port }} - {{ fault8021x }}	Fix the 802.1x error
Interface	{{ element_port }} - {{ config_error }}	Fix the configuration

Dict key	Severity	Type	Description
AMP0073	ERROR	alarm	Fabric Device(s) Communications Down

Element type	Error text	Corrective action
Switch	{{ element_id }}: Device Controller is powered off	Power on the device controller when ready to resume normal operations
Switch,ONT	{{ element_id }} is out of communication	Can be caused by a device failure, physical link failures between the device and the Verity platform or a system process going through recovery. When dependent conditions are cleared, the communications will automatically recover. If all other recovery methods have failed, reboot the device.
Switch	ACS event polling failed	This is normally caused by a transient condition due to SDLC or ACS restart. Wait for the SDLC/ACS to return to service.
Switch	Unknown Device Controller connectivity for device {{ element_id }}: {{ controller_connectivity }}	Internal System Processing Error. Create system snapshot and contact support.
Switch	Cannot configure Device Controller without controller_connectivity on {{ element_id }}!	Internal System Processing Error. Create system snapshot and contact support.

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0092	CRITICAL	alarm	ACS Down

Element type	Error text	Corrective action
ACS	ACS connection lost	Check that the SDLC and ACS are running. Contact support if the condition persists

Dict key	Severity	Type	Description
AMP0162	CRITICAL	alarm	Loss of LAG Redundancy

Element type	Error text	Corrective action
LAG	{{ element_id }}: {{ status }}	Make sure all ports are fully operational

Dict key	Severity	Type	Description
AMP0163	ERROR	alarm	Loss of LAG Redundancy

Element type	Error text	Corrective action
Olag Channel	{{ element_olag_chan }}: {{ status }}	Make sure all Ethernet ports are fully operational. Review traffic graphs for any packet layer errors. Run ping test between Router ID IP Addresses. Disable/Re-Enable ports. Reboot each switch individually
LAG	{{ element_id }}: {{ status }}	Make sure all ports are fully operational
LAG	{{ element_fabric_lagg_type }}: {{ status }}	Make sure all ports are fully operational

Dict key	Severity	Type	Description
AMP0164	WARNING	alarm	Loss of LAG Redundancy

Element type	Error text	Corrective action
LAG	{{ element_id }} is experiencing loss of redundancy	Make sure all ports are fully operational
LAG	{{ element_id }} is IDL/IDS and is out-of-service	Make sure all ports are fully operational
Switch Pair	{{ element_pair }}: {{ state }}	Make sure each Switch is fully operational
Switch	{{ element_id }}: {{ state }}	Make sure parents are fully operational

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0173	ERROR	alarm	LAG Setting Mismatch

Element type	Error text	Corrective action
LAG	{{ element_id }} - Mismatched SFP speeds in a single LAG. Failure to correct will impact LAG operation	Make sure all SFP speeds are the same
	{{ device_controller_id }} is located by LAG, but {{ element_id }} is missing provisioned LAG on port {{ element_port }}	Add Fabric LAG provisioning to the parent switch connecting to the target device.

Dict key	Severity	Type	Description
AMP0183	ERROR	alarm	Topology Detection Conflict

Element type	Error text	Corrective action
Device Controller	{{ element_id }} is a statically located switch controller but has {{ device_id }} plugged in.	Check the network fabric connections to make sure that the provisioning of the switch controller matches the desired network physical connectivity. Correction required either in provisioning or cabling.
Interface	{{ element_port }} has a connection but is provisioned as {{ eth_id }}	Remove the provisioning or move the connection
Switch	{{ element_id }} is connected to itself	Remove self connection
Interface	Unequal crosslink provisioning for connected ports: {{ element_port }} {{ conn_port }}	Make sure the provisioning matches
Switch	{{ element_id }} is treated as Edge-Device but, has a downstream connection	Remove the connection or un mark as Edge-Device
Switch	Failed to create Channel groups on {{ element_id }}	Internal System Processing Error. Create system snapshot and contact support.
Switch,ONT	Loop detected in topology on {{ element_id }} {{ element_port }}	Remove one of the connections causing the loop and determine what the intended connections should be and re-connect.
Switch,ONT	{{ element_id }} {{ element_port }} - connection conflict - {{ reject_reason }}	Remove the connection causing the issue or re-adjust the roles of the switches in the fabric.
Switch,ONT	Topo_connections - {{ topo_connections }} and VNETC - {{ vnetc }} versions do not match	This is an expected transient condition during system upgrades. If the problem persists, create a system snapshot and notify support.
Switch,ONT	{{ element_id }} {{ element_port }} - rejected connection reason - {{ reject_reason }}	Remove the connection causing the issue or re-adjust the roles of the switches in the fabric.

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0193	ERROR	alarm	Device Upgrade Failure

Element type	Error text	Corrective action
ONT	{{ element_id }} - {{ download_error }}	The system automatically retries the upgrade after a wait period.
ONT	{{ element_id }} - {{ sw_error }}	Ensure that the system has the correct core, image or firmware packages installed
Switch	Switch connection for {{ dev_id }} failed	This is normally caused by a transient condition due to switch restarting during an upgrade. Wait for the switch to return to service.
Switch	Protocol failed for dev_id {{ dev_id }}, treating as lost connection -- {{ detailed_error }}	This is normally caused by a transient condition due to switch restarting during an upgrade. Wait for the switch to return to service.
Switch	Dev_id {{ dev_id }} dl_id {{ dl_id }} upgrade failed -- {{ error_code }} {{ error_text }}	System will retry upgrade periodically after failures. If the device still fails to download make sure that the proper package is selected for the device.
Switch	Download failed with code --	System will retry upgrade periodically after failures. If the device still fails to download make sure that the proper package is selected for the device.
Switch,ONT	vNetC upgrade failure: {{ detailed_error_text }}	Internal System Processing Error. Create system snapshot and contact support.
Switch	Switch upgrade for {{ dev_id }} failed -- {{ detailed_error }}	
Switch	ONIE failure for switch {{ switch_sn }} from {{ ip_address }}: {{ detailed_error }}	Ensure switch model is supported and the necessary firmware package is loaded

Dict key	Severity	Type	Description
AMP0203	ERROR	alarm	Device Hardware Failure

Element type	Error text	Corrective action
Switch,ONT	Stats reporting is disabled	Statistics reporting has been disabled in a prior release. Contact support to have it restored
Switch,ONT	Database connection failed	Internal System Processing Error. Create system snapshot and contact support.

Dict key	Severity	Type	Description
AMP0204	WARNING	alarm	Device Hardware Failure

Element type	Error text	Corrective action
Switch	{{ element_id }} PSU failures: {{ psu_status }}	Replace or repair the PSU

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0213	ERROR	alarm	Device Internal Failure

Element type	Error text	Corrective action
Switch,ONT	Processing overran monitoring period by {{ time }}	This is normally caused by a transient condition due to system load. Contact support if the condition persists.
Switch,ONT	Processing almost overran monitoring period by {{ time }}	This is normally caused by a transient condition due to system load. Contact support if the condition persists.
Switch,ONT	Load has fallen below {{ thresh }} to {{ load }}	This indicates system load has recovered from a prior high-load condition. No action is required.
Switch,ONT	Load has risen above {{ thresh }} to {{ load }}	This indicates system load has entered a high-load condition. Contact support if the condition persists.
Switch,ONT	{{ prev_value }} is now {{ current_value }}	This message is for information only. No action is required.
Switch,ONT	{{ block_count }} postgres transaction{{ transaction_id }} blocked by {{ blocker }} blocker{{ blocker_id }} for at least {{ time }}	Consider restarting the vNetC if the condition persists. If it recurs, contact support.
Switch,ONT	Postgres notification queue utilization {{ usage }} exceeds limit {{ limit }}	This indicates system load has entered a high-load condition. Contact support if the condition persists.

Dict key	Severity	Type	Description
AMP0235	NOTICE	alarm	Device Manually Placed Out of Service

Element type	Error text	Corrective action
Switch,ONT	{{ element_id }} is marked manually out of service	Mark in service when ready to resume normal operations

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0243	ERROR	alarm	System Licensing Error

Element type	Error text	Corrective action
Switch,ONT	License not found, no provisioning attempted	Contact support to obtain a valid license for this site.
Switch,ONT	License file missing or has no licenses defined	Contact support to obtain a valid license for this site.
Switch,ONT	License file has a license entry with no tenant_id	Contact support to obtain a valid license for this site.
Switch,ONT	Site license could not be loaded even without checks -- {{ detailed_error }}	Contact support to obtain a valid license for this site.
Switch,ONT	Site license has expired	Contact support to obtain a valid license for this site.
Switch	Called with model {{ element_id }}, expecting ipho.ivas, assuming licensed	Internal System Processing Error. Create system snapshot and contact support.
Switch	License check already run, skipping additional check for device {{ element_id }}	Contact support to obtain a valid license for this site.
Switch	Using FAKE license data	Contact support to obtain a valid license for this site.
Switch	License has duplicate model	Contact support to obtain a valid license for this site.
Switch	Setting missing license tier {{ tier }} to blocked for device {{ element_id }}	Contact support to obtain a valid license for this site.
Switch	POST for {{ element_id }} does not include an 'hw_model', assuming unlicensed	Contact support to obtain a valid license for this site.
Switch	Controller {{ obj_id }} has dev_id {{ dev_id }} but is unlicensed -- checking licensing	Contact support to obtain a valid license for this site.
Switch	Could not determine license tier for model {{ model }}, controller {{ ctrl }} for dev_id {{ dev_id }}	Contact support to obtain a valid license for this site.
Switch	Controller {{ ctrl }}, dev_id {{ dev_id }} has unknown license tier {{ tier }}, marking unlicensed	Contact support to obtain a valid license for this site.
Switch	Controller {{ ctrl }}, dev_id {{ dev_id }} licensing for tier {{ tier }} is blocked, marking unlicensed	Contact support to obtain a valid license for this site.
Switch	Controller {{ ctrl }}, dev_id {{ dev_id }} licensing for tier {{ tier }} would exceed limit of {{ limit }}	Contact support to obtain a valid license for this site.
Switch	Model {{ element_id }} matched {{ n_model }} models though OUI {{ oui }} matched none, using model_id {{ model_id }}	Contact support to obtain a valid license for this site.
Switch	Model {{ element_id }} matched {{ n_model }} models, using model_id {{ model_id }} with no OUI available	Contact support to obtain a valid license for this site.
	{{ license_type }} license limit {{ limit }} - exceeded by {{ exceeded }}	Contact support to obtain a valid license
	License will expire on {{ expiration_day }}	Contact support to obtain a new valid license

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0273	ERROR	alarm	Pair Degraded

Element type	Error text	Corrective action
Switch Pair	{{ element_pair }} is out of sync	Check status of the LAG between the pair of switches and review traffic graphs for any packet layer errors. Run ping test between Router ID IP Addresses. Disable/Re-Enable ports. Reboot each switch individually

Dict key	Severity	Type	Description
AMP0305	NOTICE	incident	User Account Updated

Element type	Error text	Corrective action
Provisioning Object	{{ obj_name }} updated	This message is for informational purposes.

Dict key	Severity	Type	Description
AMP0315	NOTICE	incident	Device Created

Element type	Error text	Corrective action
Switch,ONT	{{ purpose }}Created {{ model_name }}{{ obj_name }} {{ obj_id }}	This message is for informational purposes.
Switch,ONT	{{ purpose }}{{ user_id }} created {{ model_name }}{{ obj_name }}	This message is for informational purposes.
	Bizd will run on site {{ site_id }} impacted devices {{ device_list }}	This message is for informational purposes.

Dict key	Severity	Type	Description
AMP0325	NOTICE	incident	Device Deleted

Element type	Error text	Corrective action
Switch,ONT	{{ purpose }}Destroyed {{ model_name }}{{ obj_name }} {{ obj_id }}	This message is for informational purposes.
Switch,ONT	{{ purpose }}{{ user_id }} destroying {{ element_id }}	This message is for informational purposes.

Dict key	Severity	Type	Description
AMP0345	NOTICE	incident	Object Created

Element type	Error text	Corrective action
Provisioning Object	{{ purpose }}{{ user_id }} created {{ element_id }}	This message is for informational purposes.

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0355	NOTICE	incident	Object Deleted

Element type	Error text	Corrective action
Provisioning Object	{{ purpose }}{{ user_id }} deleting {{ element_id }}	This message is for informational purposes.
Provisioning Object	{{ user_id }} failed to delete {{ element_id }}	Internal System Processing Error. Create system snapshot and contact support.

Dict key	Severity	Type	Description
AMP0365	NOTICE	incident	Object Modified

Element type	Error text	Corrective action
Provisioning Object	{{ purpose }}{{ user_id }} updated {{ element_id }}	This message is for informational purposes.

Dict key	Severity	Type	Description
AMP0395	NOTICE	incident	System incident

Element type	Error text	Corrective action
Switch Pair	{{ user_id }} completed installation, updating feature flags	
	Forget connection on {{ element_port }} --> {{ target_port }} triggered by {{ user_id }}	User manually removed automatically discovered fabric connection. No action required.
	{{ event }}	Information only. No action required.
	{{ alarm_key }} mask {{ element_pair }} {{ element_fabric_lagg_type }}{{ element_port }} expired	This message is for informational purposes.

Dict key	Severity	Type	Description
AMP0403	ERROR	alarm	Failed login attempt

Element type	Error text	Corrective action
	Failed login attempt by IP: {{ ip_address }}, username: {{ username }}	This message is for informational purposes.
	Failed connection attempt by IP: {{ ip_address }}, common name: {{ cn }}, message: {{ message }}	Check correct CRLs and Certificate Chains are loaded

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0413	ERROR	alarm	Monitoring Down

Element type	Error text	Corrective action
Switch	Monitoring VM not reachable, monitoring feature not available.	Ensure that the Monitoring VM is properly installed and can communicate with the vNetC.

Dict key	Severity	Type	Description
AMP0425	NOTICE	incident	Name Changed

Element type	Error text	Corrective action
Switchpoint	From {{ old_name }} to {{ new_name }}	This message is for informational purposes.

Dict key	Severity	Type	Description
AMP0433	ERROR	alarm	System Component Provisioning Failure

Element type	Error text	Corrective action
	{{ element_id }} - {{ detailed_error }}	Remove user provisioning or add redundancy

Dict key	Severity	Type	Description
AMP0443	ERROR	alarm	Failed login attempt in last 24 hours

Element type	Error text	Corrective action
	Failed login attempt (check history for details)	This message is for informational purposes.

Dict key	Severity	Type	Description
AMP0453	ERROR	alarm	Failed login attempt by blocked IP address

Element type	Error text	Corrective action
	Failed login attempt by blocked IP address: {{ ip_address }}	This message is for informational purposes.

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0463	ERROR	alarm	BGP Issue

Element type	Error text	Corrective action
	{{ element_id }} <--> {{ sub_id }} : {{ status }}	Ensure physical connections are in service with no errors. Ensure that the path and the devices associated with the BGP neighbor are all in service. Re-check the BGP parameters in the Gateway object and Gateway Port Profile, and that they match the external neighbor

Dict key	Severity	Type	Description
AMP0464	WARNING	alarm	BGP Issue

Element type	Error text	Corrective action
Switch	{{ element_id }} <--> {{ sub_id }}: {{ status }}	Ensure physical connections are in service with no errors. Ensure that the path and the devices associated with the BGP neighbor are all in service. Re-check the BGP parameters in the Gateway object and Gateway Port Profile, and that they match the external neighbor

Dict key	Severity	Type	Description
AMP0475	NOTICE	alarm	Read Only Pending

Element type	Error text	Corrective action
Switch	{{ element_id }} is read only with pending changes	Review reasons for pending changes. This can be caused by direct provisioning on the device that conflicts with the Verity intended configurations, or changes as a result of a software upgrade. Also can be caused by changes made to the system provisioning while the device was in read only mode. When ready to resume normal operations, take out of read only mode to apply changes, or revert provisioning changes manually on the device.

Dict key	Severity	Type	Description
AMP0485	NOTICE	incident	Global read only modified

Element type	Error text	Corrective action
System	{{ user_id }} set global read only to {{ global_read_only }}	This message is for informational purposes.

Dict key	Severity	Type	Description
AMP0494	WARNING	alarm	Onboarding enabled

Element type	Error text	Corrective action
System	{{ user_id }} enabled onboarding with stop time {{ stop_time }}	This message is for informational purposes.

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0505	NOTICE	incident	Onboarding stopped

Element type	Error text	Corrective action
System	stopped by {{ user_id }}	This message is for informational purposes.

Dict key	Severity	Type	Description
AMP0515	NOTICE	alarm	Unsupported Feature

Element type	Error text	Corrective action
System	Unsupported Capabilities present on the system, see report for more details	Certain provisioned features are not supported by Verity and/or the physical device. Feature capabilities are display at the top of the switch in the GUI. A warning triangle on the switch indicated which specific feature is causing the alarm. This error is a warning to the user and can be ignored, or the provisioning objects can be modified to eliminate the unsupported feature.
Interface	{{ element_port }} - The chosen breakout override is not supported	Update the Port Breakout selection

Dict key	Severity	Type	Description
AMP0525	NOTICE	alarm	Unregistered Device Present

Element type	Error text	Corrective action
System	Unregistered device present on the system, see report for more details	Usually a transitory state on the system as new devices are discovered. Allow the system initialization of the device to complete.

Dict key	Severity	Type	Description
AMP0535	NOTICE	incident	Global firmware upgrades modified

Element type	Error text	Corrective action
System	{{ user_id }} {{ state }} global firmware upgrades	This message is for informational purposes.

Dict key	Severity	Type	Description
AMP0555	NOTICE	incident	Import incident

Element type	Error text	Corrective action
Provisioning Object	Completed {{ user_id }}'s import, adds: {{ adds }}, deletes: {{ deletes }}, changes: {{ changes }}	This provides a short summary of what changed during a user import.
Provisioning Object	{{ user_id }} started import	This indicates that a user started an import, either through the GUI or the NBI

Copyright BE Networks 2026

Dict key	Severity	Type	Description
AMP0564	WARNING	alarm	External backup failed

Element type	Error text	Corrective action
	Backup to external server {{ server }} failed: {{ message }}	Ensure external server is accessible and necessary authentication keys are installed on it

Dict key	Severity	Type	Description
AMP0575	NOTICE	incident	User action

Element type	Error text	Corrective action
Switch,ONT	{{ user_id }} rebooted {{ element_id }}	Indicates that a user started a reboot of a device.

Dict key	Severity	Type	Description
AMP0585	NOTICE	incident	NBI incident

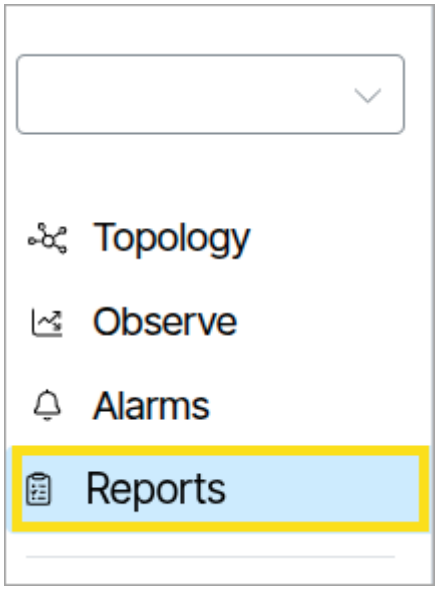
Element type	Error text	Corrective action
Provisioning Object	Completed NBI request by {{ user_id }}, adds: {{ adds }}, deletes: {{ deletes }}, changes: {{ changes }}	This provides a short summary of what changed during a NBI user request.

Copyright BE Networks 2026

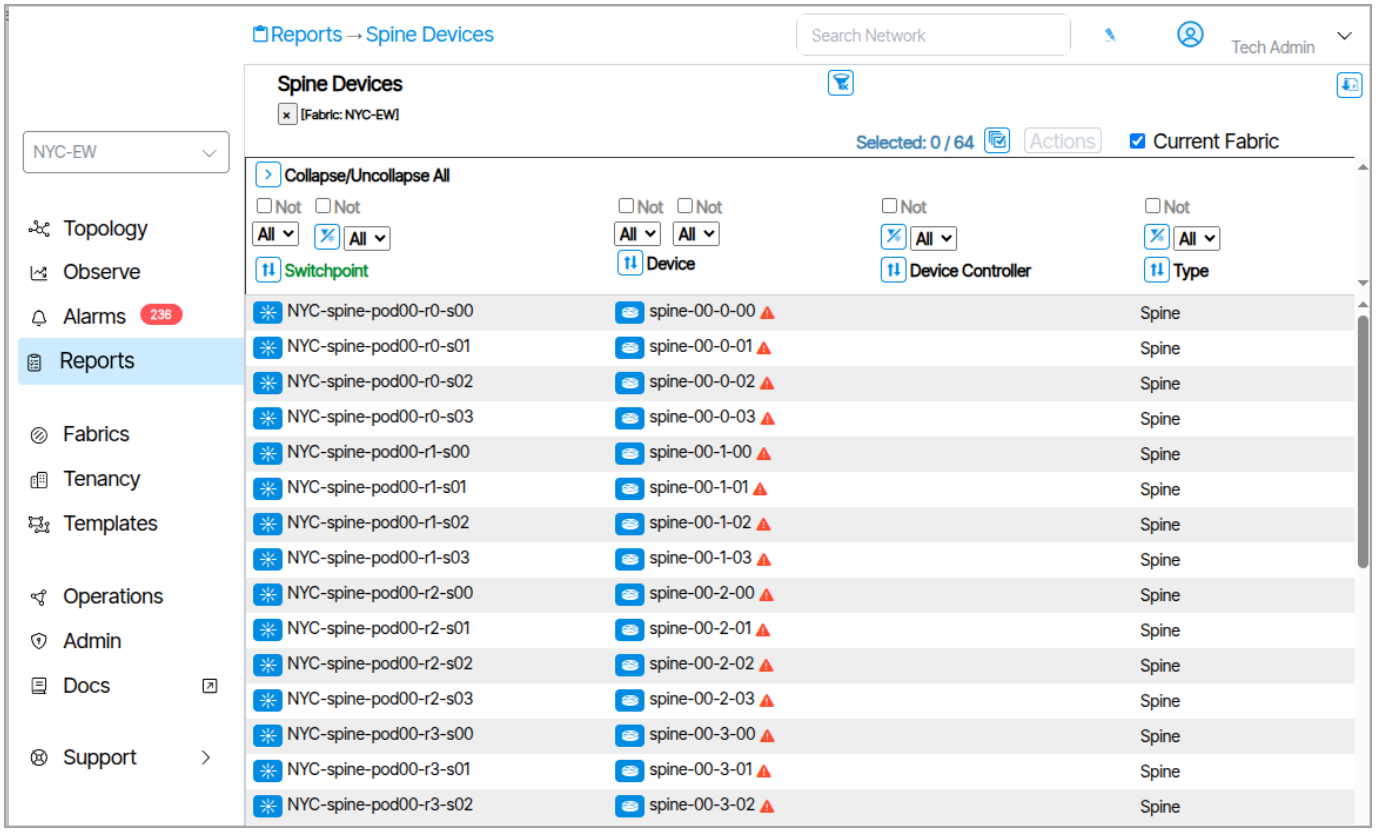
Reports

4.6.7 Reports

The **Reports** section presents an organized overview of all available *Reports*. Reports are tabular views of provisioned objects and detected elements within the network.



Selecting a Report opens a text-based list that can be sorted, highlighted, exported to CSV, and navigated from. In the example below, the Report titled **Spine Devices** is opened.



Reports can be filtered and customized to create individual operator-specific reporting on any element in Verity.

Copyright BE Networks 2026

Filtering and Sorting

FILTER COLUMN ON OBJECT TYPE

Filter objects using **Filter Column on Object Type** at the top of each column.



REPORT FILTERING USING REGEX

Click the icon shown below to apply regular expressions to report filters. Type the desired Regex command in the field that appears.



SORT COLUMNS

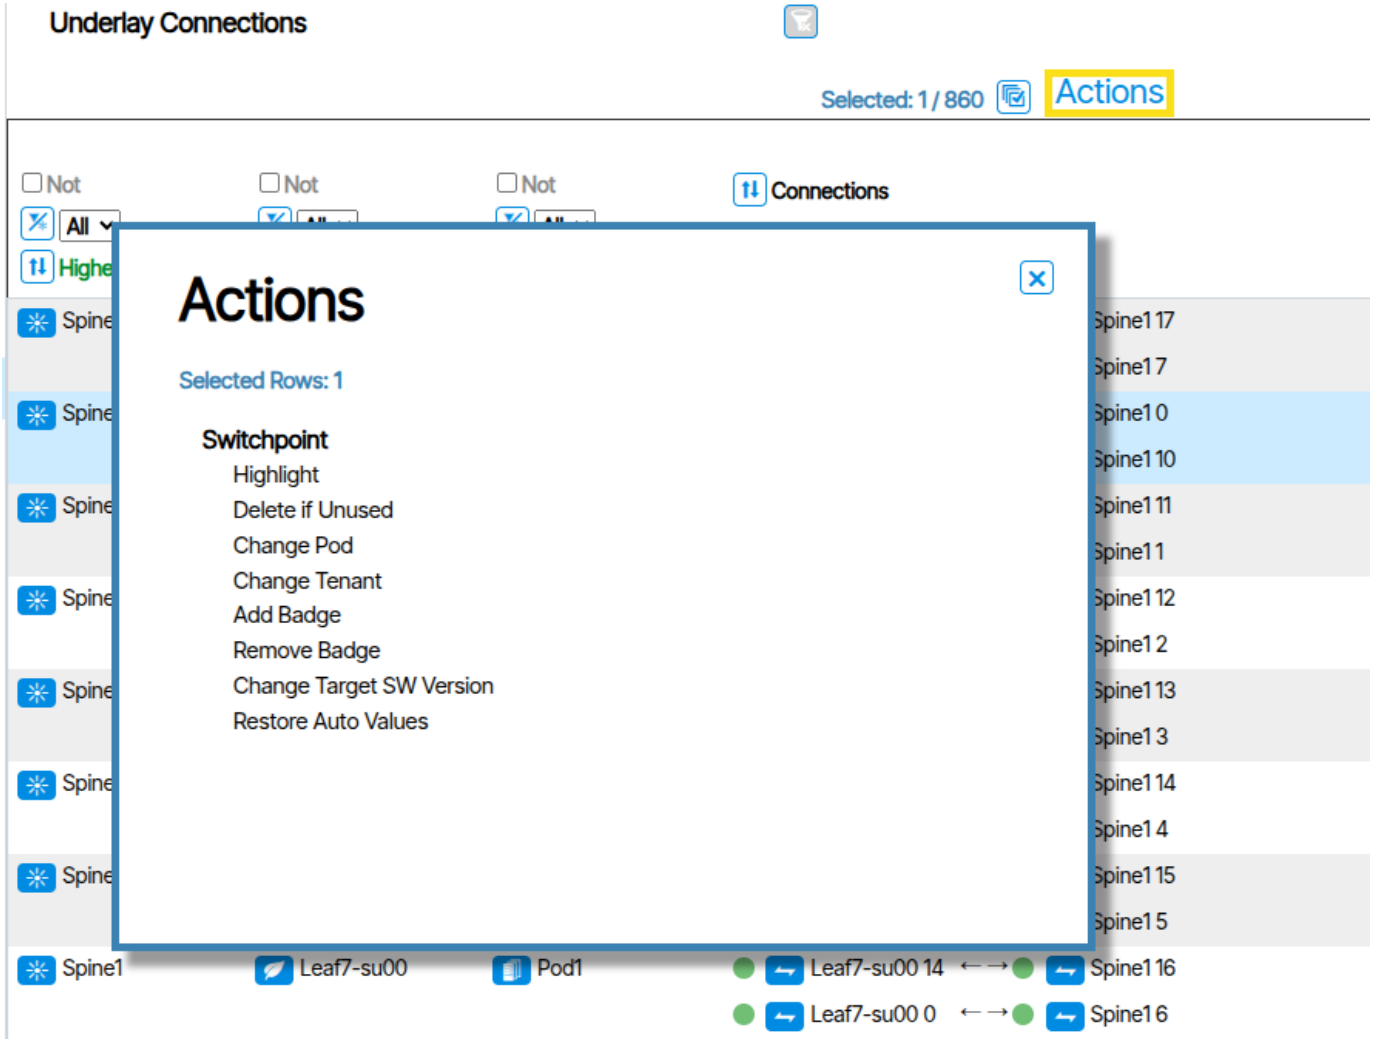
Sort columns using the column sort tool.



Report Actions

To apply actions to **Report** items, select the item(s) you want to affect, click the **Actions** button, and choose the desired action. All selected Report items are affected by the change.

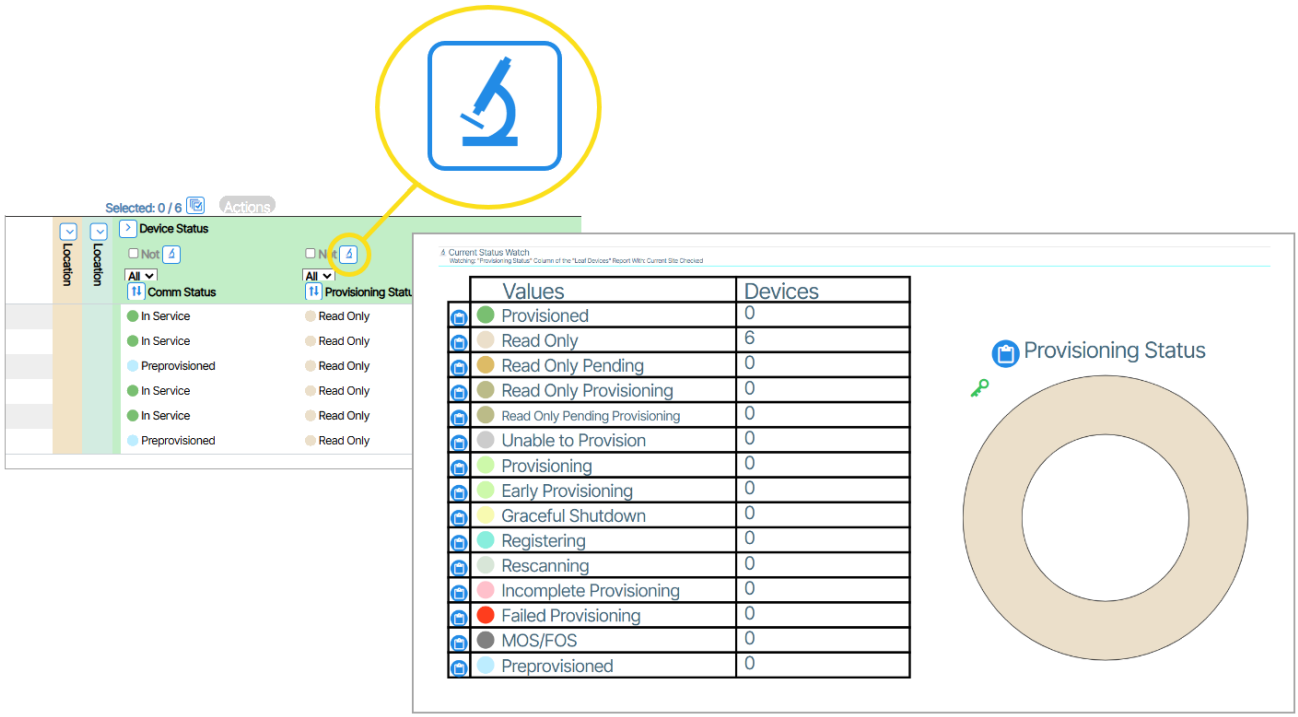
Copyright BE Networks 2026



REPORT SLICE VIEWS

Report Slice View displays a pie graph view of the column in the report. To view the graph click the slice button

Copyright BE Networks 2026



Copyright BE Networks 2026

Layer-3

Report Name	Content
Fabrics/LAN	A Fabric is an isolated, self-contained arrangement of network interconnections—switches, links, and logical structure—that forms a network topology. The Fabrics/LAN Report is a tabular view of values, settings and information for all Fabrics and LAN topologies.
Fabrics Collections	Fabrics are organized into tree-node containers called Fabric Collections. The Fabrics Collections Report is a tabular view of values, settings and information for all Fabric Collections.
Tenant Route Tables	Tenant Route Table displays the virtualized routing table (VRF) that's specific to a single tenant in a shared infrastructure.
Underlay Connections	An Underlay Connection is the physical network infrastructure that provides the foundation for packet transport between devices. The Underlay Connection Report provides a list of physical connections used for inter-switch Layer-2 & Layer-3 connections.
External BGP Connection	An external BGP (eBGP) connection is a Border Gateway Protocol link established between network routers. It allows separate organizations or network domains to exchange routing information and connect to each other while maintaining their independent routing policies and configurations. The External BGP Connection Report provides a list of the configured BGP neighbors as well as their current state.
DHCP Lease	A DHCP Lease is a temporary assignment of an IP address and related network configuration parameters to a device by a DHCP server. The DHCP Report is a tabular list of all DHCP leases provided by Verity on the management network.
DHCP Assigned IPs	DHCP Assigned IPs are IP addresses that have been dynamically allocated to devices on a network by the DHCP server. The DHCP Assigned IPs Report lists all of the IP addresses assigned by the Verity server.
DHCP Relays	A DHCP Relay is a service running on a network device (typically a router) that forwards DHCP requests from clients to a DHCP server located in a different network segment or subnet. The DHCP Relay Report lists all of the configured DHCP relays in the managed network.
Internal IPs	Internal IPs are private IP addresses assigned to devices within a local network that are not directly accessible from the public internet.
Gateways	A Gateway is an object that defines an external BGP or static route connection for a Tenant or the Underlay fabric. The Gateways Report is a tabular list of all configured Gateways and their current state.

Copyright BE Networks 2026

Layer-2

Report Name	Content
Isolated LANs	An isolated LAN is a local area network that is intentionally disconnected from any other network, including the internet and any other internal networks. This means that devices on the isolated LAN can only communicate with each other, and cannot communicate with devices outside of that LAN.
Underlay Route Tables	An underlay route table is used by physical network devices to determine the best path for forwarding traffic within the underlying physical network infrastructure. Routing protocols like OSPF, IS-IS, and BGP populate these tables to ensure all devices in the underlay can reach each other, providing essential connectivity for overlay networks.
Active Captures	Refers to the process of intercepting and recording network traffic as it occurs. The Active Captures Report is a tabular list of actions, values, and settings for Capture processes that are still ongoing.
LLDP on Edge Ports	LLDP is a protocol that allows network devices to advertise information like their identity and capabilities to neighboring devices on a local network. On edge ports, LLDP can be used to auto-configure VoIP phones, inventory devices, and troubleshoot connectivity, but it may be disabled for security reasons to prevent unauthorized disclosure of network topology.
MACs	The **MAC Address Workbench** (MACs) lets you search device by mac address. Clicking the **Search** button without specifying a mac address displays all devices (). From within a list of devices, clicking () opens a new tab to search on the corresponding mac address ()
MAC Export	MAC Export provides a summary MAC report for all devices connected to the network's user facing ethernet ports. Administrators can automate this tool to generate a daily report or use it on a case-by-case basis as needed. To manually run a report, enable the checkbox for all services you want to include, click** Run Report Now**, and wait for the **Download Latest Report** icon to turn blue. Once it does, click the icon to download the CSV file.
LAGs	A LAG (Link Aggregation Group) is a collection of Ethernet ports that are grouped together and behave as a single logical interface, providing increased bandwidth, redundancy, and resilience. The LAGs Report is a tabular list of actions, values, and settings for all LAGs.
Services	A Service is a Layer-2 broadcast domain (VLAN) that can optionally be routed through the Layer-3 fabric as VXLAN frames, allowing devices connected to different parts of the network to appear as if they are on the same local segment. The Services Report is a tabular list of actions, values, and settings for all Services.

Copyright BE Networks 2026

Layer-1

Report Name	Content
Ports	The Ports Report is a tabular list of actions, values, devices and settings for all Ports.
Fabric	The Fabric Report is a tabular list of actions, values, devices and settings for all Fabric Ports.
Provisioned	The Provisioned Report is a tabular list of actions, values, devices and settings for Provisioned Ports.
Spare	The Spare Report is a tabular list of actions, values, devices and settings for Spare Ports.
Optics	In networking, optics generally refers to the components and technologies that utilize light to transmit data over fiber optic cables. Instead of electrical signals used in traditional copper-based networking, optics leverages light pulses to carry information, enabling significantly higher bandwidth and longer transmission distances with less signal degradation.
SFP	SFP stands for Small Form-factor Pluggable, which is a type of transceiver used in networking equipment. It is a hot-swappable component that can be inserted into a switch to provide connectivity.
Static Connections	A Static Connection is a connection that is not dynamic. The Static Connections Report is a tabular list of actions, values, and settings for Static Connections.
Missing Port Connections	A Missing Port Connection is a port with a connection provisioned but no matching physical link detected from the expected downstream device. The Missing Port Connections Report is a tabular list of actions, values, devices and settings for all Missing Port Connections.

Copyright BE Networks 2026

Devices

Copyright BE Networks 2026

Report Name	Content
Switchpoint	A Switchpoint represents a logical location for a device in the network and defines a connection point, such as where leaf 1 connects to spine 3, independent of the specific device installed there. The Switchpoint Report is a tabular view of values, settings and information for all Switchpoints.
Preprovisioned Switchpoint	A Preprovisioned Switchpoint is a Switchpoint that has been configured before a physical device has been connected to it, allowing operators to define a device's settings in advance of installation. The Preprovisioned Switchpoint Report is a tabular view of all Switchpoints that are awaiting a device assignment.
Switch Bundles	A Bundle is a collection of the various templates assigned to a Switchpoint. The Switch Bundles Report is a tabular view of the Bundles assigned to each switch.
Bundle CLI Commands	The Bundle CLI Commands Report is a tabular list of the CLI commands generated from the templates assigned to a switch's Bundle.
Devices	The Devices Report is a tabular view of values, settings and information for the following devices: - **Management Devices** - **Superspine Devices** - **Spine Devices** - **Leaf Devices** - **Packet Broker Devices**
Management	Management Devices are switches that provide connectivity to the out-of-band management ports of switches across other fabrics. The Management Report is a tabular view of values, settings and information for all Management Devices.
SuperSpine Groups	SuperSpine Groups are collections of Superspine devices grouped together within a fabric. The SuperSpine Groups Report is a tabular view of values, settings and information for all SuperSpine Groups.
Superspine	A Superspine is a third switch layer used to connect Pods to each other. The Superspine Report is a tabular view of values, settings and information for all Superspine Devices.
Spine	Spine devices interconnect all Leaf switches in a full-mesh topology, forming the core of a spine-leaf network. The Spine Report is a tabular view of values, settings and information for all Spine Devices.
Planes	A Spine Plane is a collection of Superspine devices that connects to a Spine within one or more PODs, forming a dedicated inter-layer fabric in multiplane architectures. The Planes Report is a tabular view of values, settings and information for all Spine Planes.
Pods	A POD is a grouping of connected Spine and Leaf switches, allowing for added network capacity to be deployed without redesigning the entire network. The Pods Report is a tabular view of values, settings and information for all PODs.
Leaf	Leaf devices are access switches that aggregate traffic from servers and connect directly into the Spine or network core. The Leaf Report is a tabular view of values, settings and information for all Leaf Devices.
Racks	A Rack represents the physical location of a device within a data center. The Racks Report is a tabular view of values, settings and information for all Racks.
Packet Broker	A Packet Broker is a network device placed between network traffic sources and monitoring tools or other destinations. The Packet Broker Report is a tabular view of values, settings and information for all Packet Broker Devices.
Servers	The Servers Report is a tabular view of values, settings and information for all Servers connected to the network.
Device Comm Status	Device Comm Status is the state a device is in and is expressed by different colors that represent different states. The Device Comm Status Report is a tabular view of each device and their current state.
Traffic Mirrors	A Traffic Mirror is a network monitoring tool that copies traffic from a source port to a destination port so it can be analyzed. The Traffic Mirrors Report is a tabular view of values, settings and information for all configured Traffic Mirrors.
Image Package Usage	Image Package Usage Report lists the currently used SDLC image by name and version number.
Pending Changes	

Copyright BE Networks 2026

Report Name	Content
	Pending Changes includes all switches that have .system.configuration.target-pending set to true (and are not read-only).

System

Report Name	Content
Device Thresholds	Thresholds allow administrators to create rules that monitor a specific metric with configurable conditions (greater than, less than, equals, etc.) and severity levels, generating an alert when a device's metric meets those conditions. The Device Thresholds Report is a tabular view of values, settings and information for all Device Thresholds.
Port Thresholds	Thresholds allow administrators to create rules that monitor a specific metric with configurable conditions (greater than, less than, equals, etc.) and severity levels, generating an alert when a port's metric meets those conditions. The Port Thresholds Report is a tabular view of values, settings and information for all Port Thresholds.
SDLCs	The SDLCs Report is a tabular view of each SDLC and their current state.
Badges	Badges serve as a grouping mechanism for endpoints within the system. They allow you to categorize endpoints based on specific criteria—such as their function, vendor (e.g., Dell), or any other organizational preference. By assigning the same badge color and number to related endpoints, you can easily group and identify them according to these classifications.
Tier Licensing Utilization	The Tier Licensing Utilization Report is a tabular list of each device's licensing status.

Copyright BE Networks 2026

Button Colors | Colors | System Colors

4.6.8 Status Codes

Device Provisioning Status

Various colors are used to represent different provisioning states of network switches. Each state is paired with a specific color code (hex value).

Copyright BE Networks 2026

Message

Device Status is displayed in **Reports/Report Slice** 

Copyright BE Networks 2026

Color	Status	Meaning
 (Green / #79BF71)	**Provisioned**	Provisioning is considered OK, with no detected pending work or overriding provisioning condition.
 (Green / #CDF9AA)	**Provisioning**	Configuration work is pending or in progress: desired and current provisioning versions differ, pending changes are detected, or the controller/ switch reports configuring.
 (Light green / #CDF9AA)	**Early Provisioning**	The GUI anticipates pending provisioning and temporarily replaces an otherwise Provisioned status before the normal status catches up.
 (Yellow / #F8FAB1)	**Graceful Shutdown**	The switch's graceful-shutdown workflow has started. The implementation selects this when the reported state is present and is not `not-started`.
 (Teal / #8EDDDD)	**Registering**	The device is registering and is not yet available; readable port definitions are not yet present.
 (Gray / #D8E6D8)	**Rescanning**	The device is being rescanned.
 (Pink / #FFC0CB)	**Incomplete Provisioning**	Required provisioning setup is missing or inconsistent—for example, a missing endpoint, a missing device bundle, an invalid topology association, or conflicting port assignments.
 (Red / #FF3D1D)	**Failed Provisioning**	A provisioning error or switch/controller failure has been reported. The tooltip provides the specific failure.
 (Grey / #CCCCCC)	**Unable to Provision**	Provisioning cannot currently proceed—for example, because communication is unavailable, synchronization or registration data is missing, or an upstream dependency is unavailable.
 (Teal / #BDEDFF)	**Preprovisioned**	Configuration exists for a placeholder/mock device before the real device is present.
 (Pale yellow / #F4F395)	**Initializing Switch**	The switch-management status reports `initial upload` or `agent started`.
 (Beige / #EBDFCA)	**Read Only**	Read-only mode prevents configuration writes to the switch; no pending-change flag is present.
 (Brown-Orange / #DDBB66)	**Read Only Pending**	Read-only mode is enabled and pending changes are recorded.
 (Light Brown / #BBBB88)	**Read Only Provisioning**	Read-only mode is enabled while the provisioning logic detects pending/in-progress work; the separate pending-change flag is absent. The label does not by itself mean writes are permitted.
 (Muted khaki / #BBBB88)	**Read Only Pending Provisioning**	The same read-only provisioning condition, with the pending-change flag also present.
 (Gray / #808080)	**Maintenance Out of Service (MOS) / Fault out of Service (FOS)**	Device is out of service.

Copyright BE Networks 2026

Topology Device Status

These colors apply without selection, highlighting, or other display overrides.

Copyright BE Networks 2026

Color	Status	Meaning
 (Near-white / #F8FFFF)	**Provisioned / normal device**	Default background for physical hardware. A Provisioned state retains this fill; the green status indicator is separate.
 (Light blue / #BDEDFF)	**Preprovisioned / No Device / mock device**	A placeholder stands in for expected hardware that is not yet present. Also used for preprovisioned endpoints.
 (Yellow-green / #DDF9AA)	**Provisioning**	Configuration work is pending or in progress.
 (Yellow-green / #DDF9AA)	**Early Provisioning**	The GUI anticipates pending provisioning before the normal provisioning status catches up.
 (Pale yellow / #F8FAB1)	**Graceful Shutdown**	The switch's graceful-shutdown workflow has started.
 (Mint green / #88EEDD)	**Registering**	Registration is in progress and normal provisioning is not yet available; readable port definitions are not yet present.
 (Muted light green / #C1D69A)	**Rescanning**	A rescan is in progress.
 (Pink / #FFC0CB)	**Incomplete Provisioning**	Required configuration is missing or inconsistent—for example, a missing endpoint or conflicting provisioning assignments.
 (Light red / #FF6666)	**Failed Provisioning**	A provisioning error or switch/controller failure has been reported.
 (Very light gray / #E8E8E8)	**Unable to Provision / Upstream Fabric Port Down**	Provisioning cannot proceed, or an upstream connection is unavailable. The tooltip identifies the cause.
 (Gray / #BBBBBB)	**MOS/FOS**	Marked out of service, or forced out of service by an upstream condition.
 (Pale yellow / #F4F395)	**Initializing Switch**	Switch-management status reports initial upload or agent startup.
 (Pale beige / #EBDFCA)	**Read Only**	Read-only mode prevents configuration writes; no pending-change flag is present.
 (Tan / #DDBB66)	**Read Only Pending**	Read-only mode is enabled and pending changes are recorded.
 (Muted khaki / #BBBB88)	**Read Only Provisioning**	Provisioning work is detected while read-only mode is enabled; the separate pending-change flag is absent.
 (Muted khaki / #BBBB88)	**Read Only Pending Provisioning**	Provisioning work is detected while read-only mode is enabled, and the pending-change flag is present.
 (Salmon pink / #FFA0A0)	**Switch without an endpoint**	A real switch is not associated with a switch endpoint. This has its own color rather than the ordinary Incomplete Provisioning pink.
 (Pink / #FFC0CB)	**Licensing restriction**	Styling for a placeholder that cannot be licensed under the applicable endpoint-tier

Copyright BE Networks 2026

Color	Status	Meaning
		check. This is distinct from Incomplete Provisioning despite sharing its color.
 (Blue / #A7DCFF)	**Unmanaged device / unmanaged external gateway**	Attached unmanaged hardware or a gateway is drawn through the discovery/link-status path.
 (Blue-gray / #DAE0E4)	**Previously discovered unmanaged device, link down**	Disabled styling is applied when discovery information exists but the link is down.

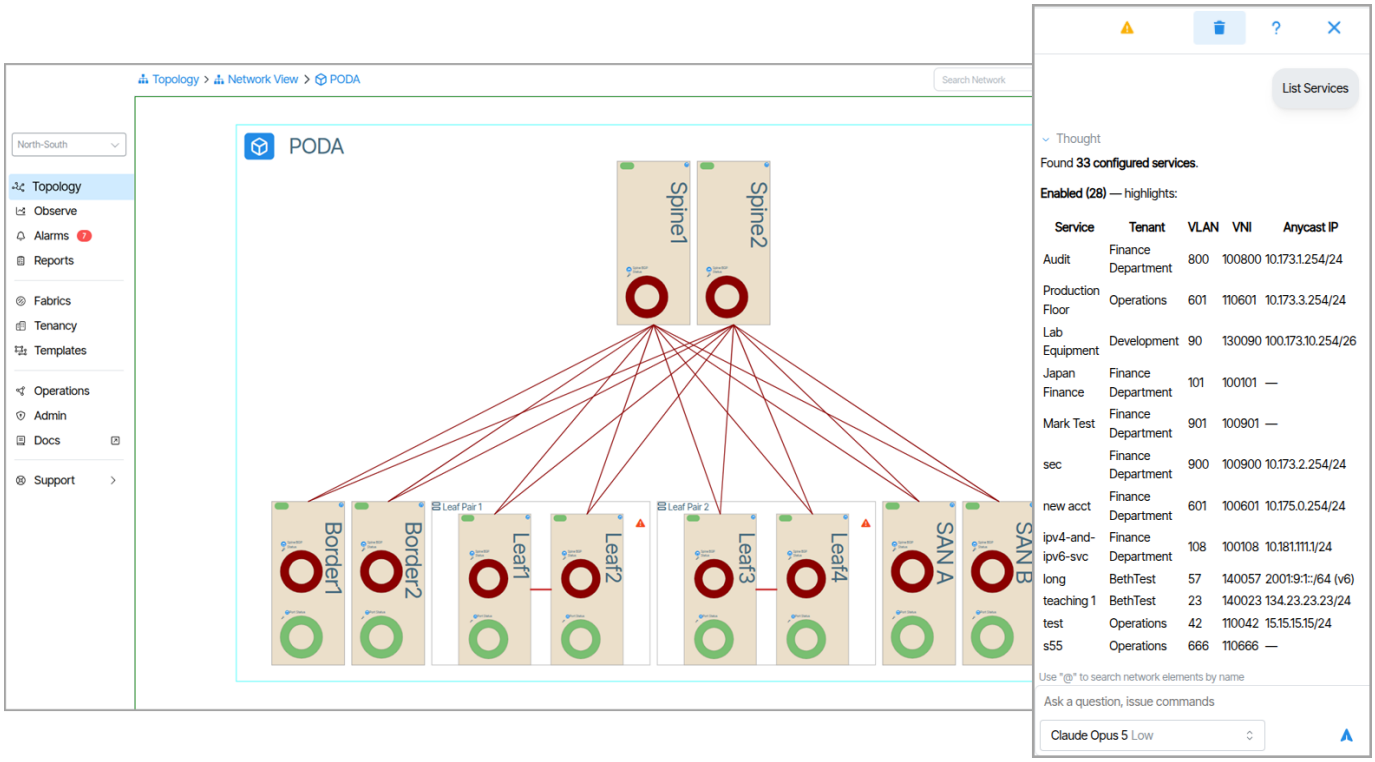
Copyright BE Networks 2026

4.6.9 AI Ops

Installation

For information on installing, disabling or updating AI Ops, see the section on AI Service System Applications.

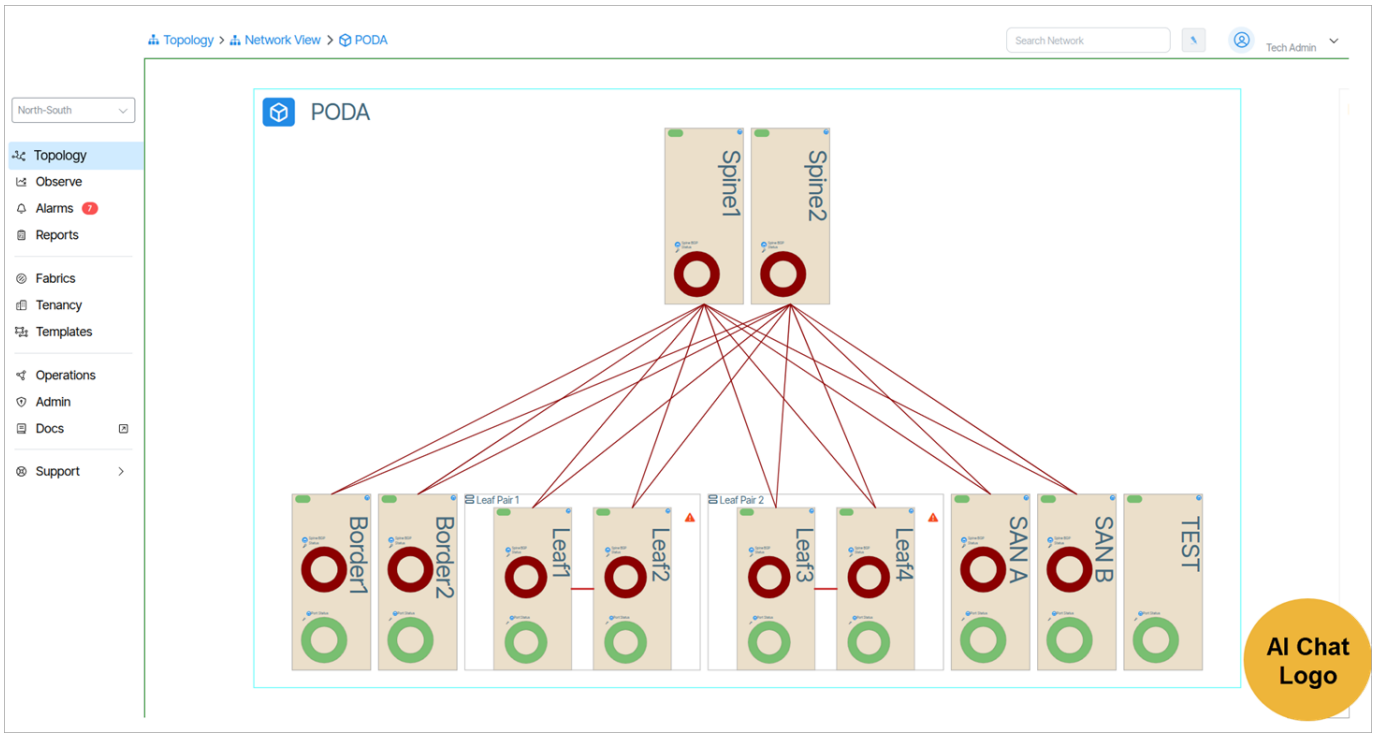
AI Ops is an AI-powered messaging assistant designed to help network professionals monitor and configure their network architecture. Through its intuitive interface, the assistant enables administrators to quickly diagnose problems and identify solutions.



Getting Started with AI Ops

Click the **AI Chat Logo** button to open AI Ops.

Copyright BE Networks 2026

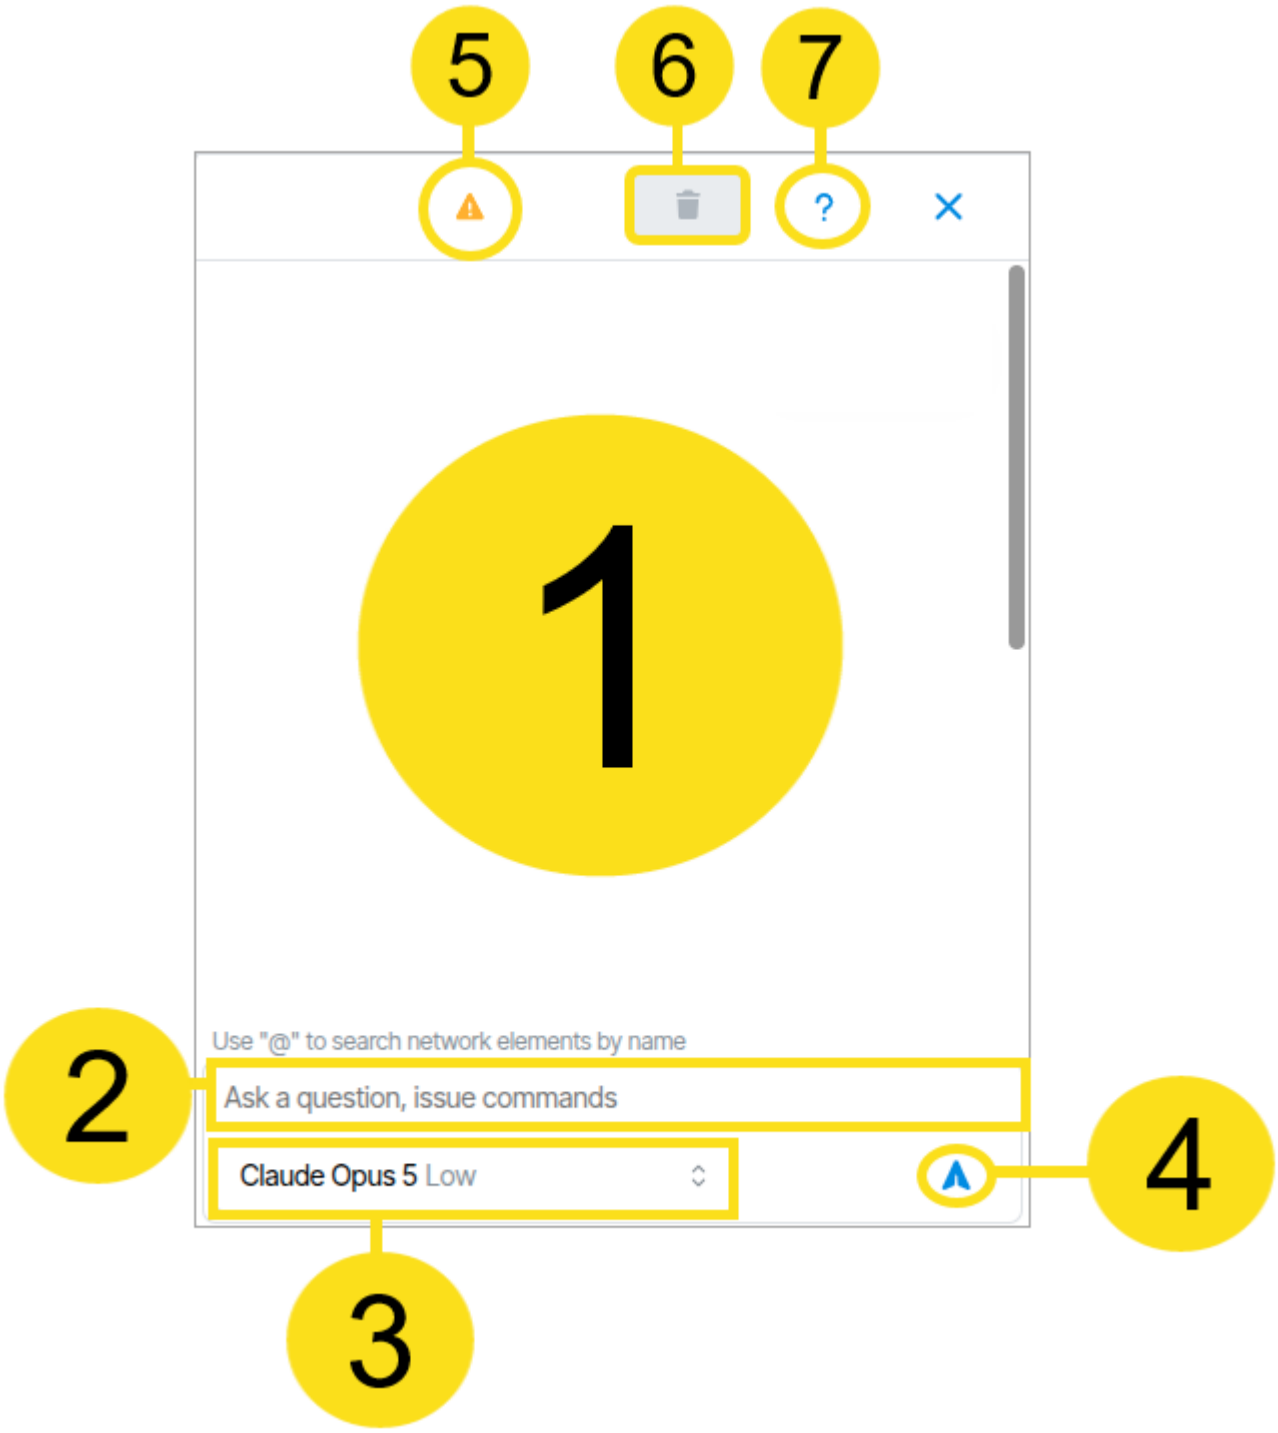


CHAT MESSAGE INTERFACE

The chat message interface is where users interact with AI Ops.

1. AI Ops Response Field
2. User Submission Prompt Field
3. LLM Selection
4. User Submission Button
5. Name Collision Warnings
6. Clear Conversation and Context button
7. Help Button

Copyright BE Networks 2026

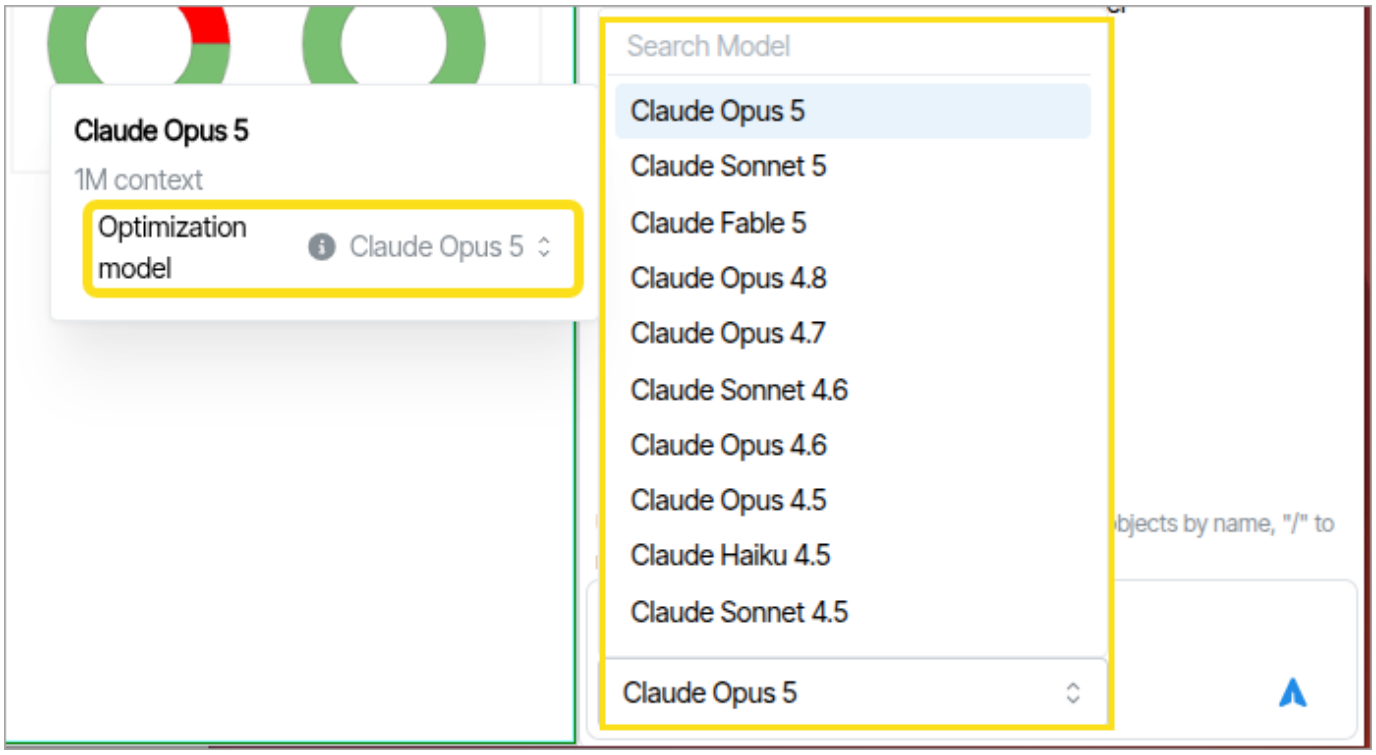


LLM MODEL SELECTION

> **Note**
>
> Official LLM support is limited to Anthropic, TogetherAI and OpenAI models.

The **Model** selection is listed as a dropdown menu at the bottom of the chat window.
When selecting an LLM, an additional **Optimization Model** may be available.

Copyright BE Networks 2026

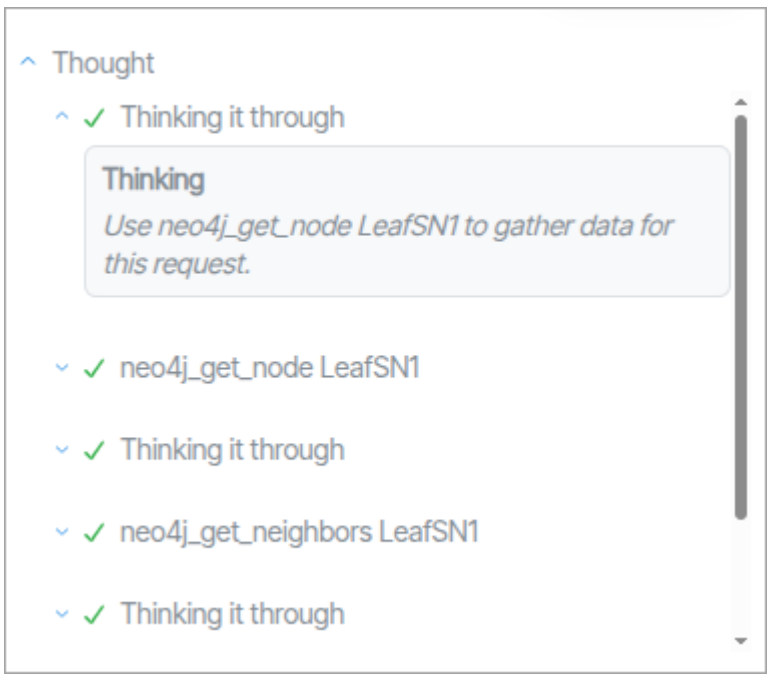


PROMPTING

When submitting prompts, the system displays a thinking indicator.



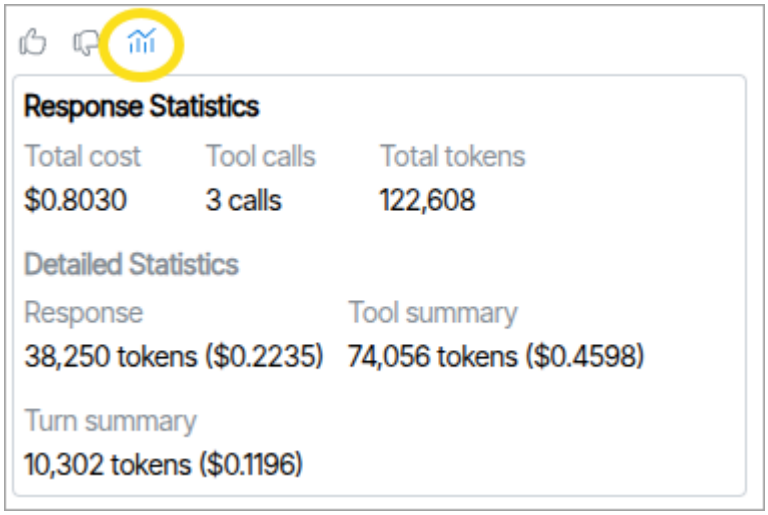
When a prompt is complete, a summary of the process is expressed by an entry titled **Thought**. This entry is composed of a list of actions taken to create the prompt response. The actions can be inspected by clicking on them.



Copyright BE Networks 2026

Response Statistics

Clicking the **Response Statistics** button at the bottom of a reply opens a window displaying statistics about the prompt response, including token count, number of tool calls, and cost. Response statistics are viewable by all users. Additional **Detailed Statistics** are available only to administrators with a permission level of Admin Tech.



Feedback

Thumbs Up/Thumbs Down

Users of AI Ops can help guide more accurate replies by clicking the thumbs-up or thumbs-down feedback icons after receiving a response.

Fuzzy Matching

AI Ops is designed to anticipate user intent when object names are queried. Since object names may occasionally differ slightly from what users type, AI Ops employs a technique known as *fuzzy matching* to interpret and correct user input. For example, if a user searches for an object named "Spine-1" but enters "Spine 1" (omitting the hyphen), AI Ops will attempt to identify and match the intended object. AI Ops notifies the user of the change by posting the word "corrected" beneath the query.

If the current query results in multiple potential matches—creating naming collisions—fuzzy matching is temporarily disabled. In such cases, AI Ops prompts the user to clarify their query.

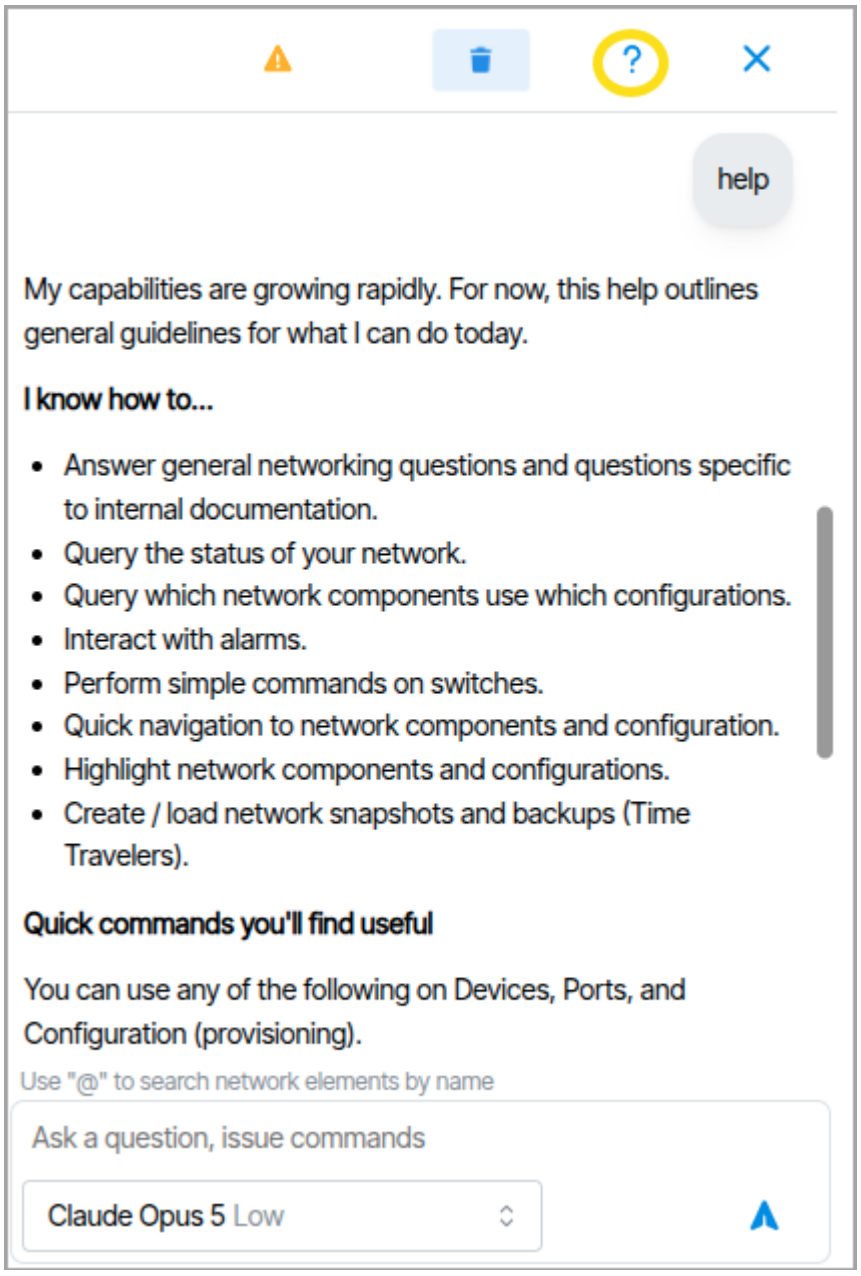
HELP

AI Ops provides a useful help menu with examples of its capabilities. To view **Help**, click the question mark icon in the upper right corner of the AI Ops chat window.

Note

AI Ops scope is limited to specific objects and tasks. This documentation provides a general overview of its capabilities. For a detailed list of applicable objects and actions please view the **Help** option in the application.

Copyright BE Networks 2026



Commands

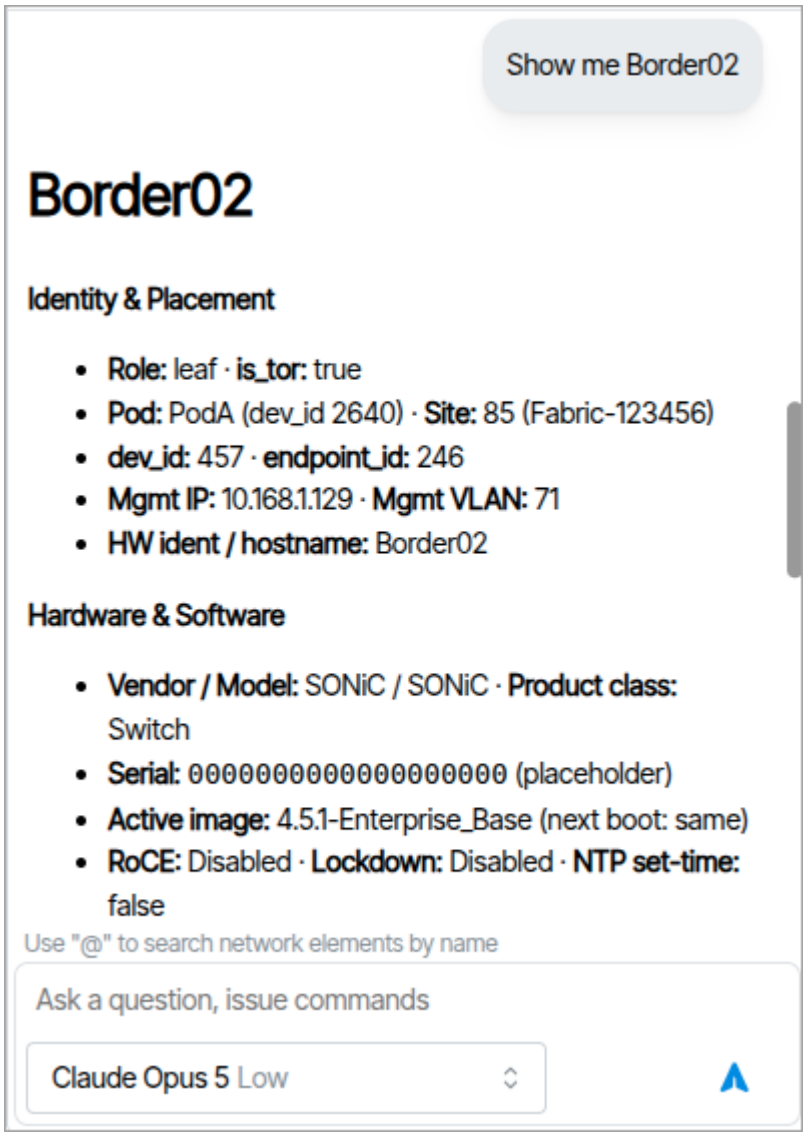
AI Ops responds to user submitted prompts by displaying data in various formats and visual styles. In addition to *displaying* data, AI Ops also responds to user queries by *executing site-wide commands* that impact overall system operations. A full list of commands is available in the help section.

Displaying Objects

AI Ops displays objects that match the query scope.

The query in the example below returns Border02 device information.

Copyright BE Networks 2026



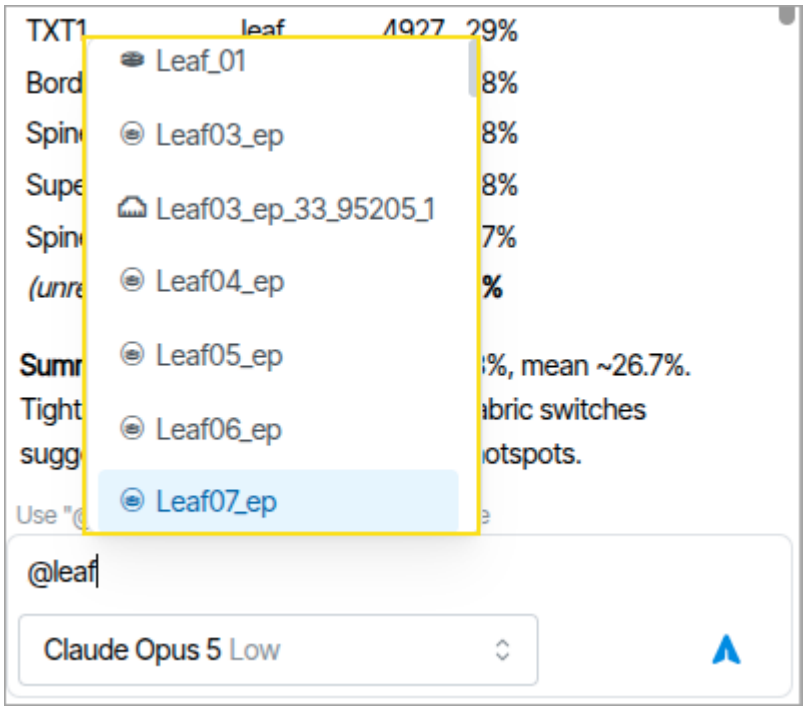
Verity Documentation Reference

When prompting AI Ops, material from the official documentation may be presented and/or requested by the user. This information is usually presented in aggregate form from different sections of the documentation, with individual reference links to the source material included. Responses may include documentation content as part of a larger reply.

Instant Search

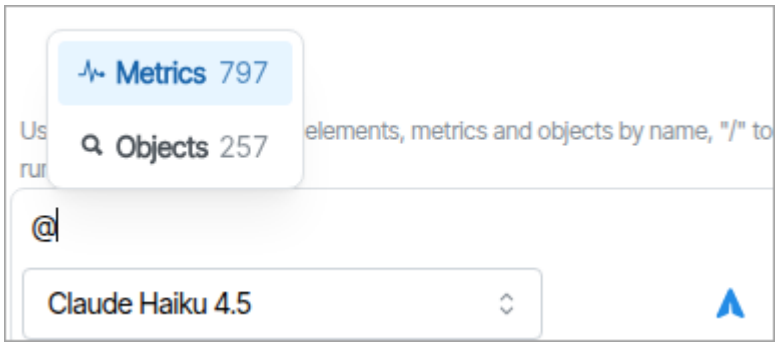
The instant search allows users to quickly find objects by name. To use this feature, preface the search with the @ symbol and type the name of the object or device you are searching for. Use the up and down arrows to choose from the presented items, and press **Tab** to auto-populate the selection in the submission field.

Copyright BE Networks 2026



Instant Search for Metrics and Objects

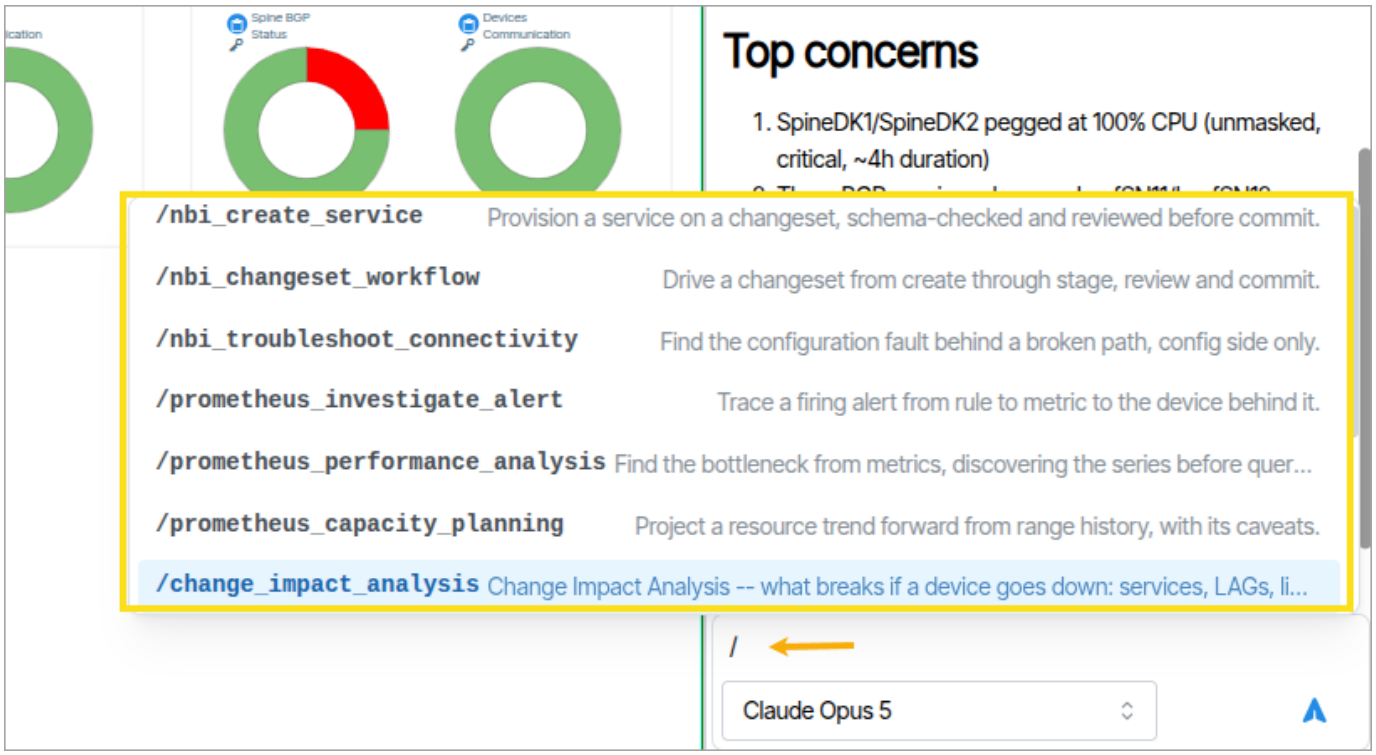
When typing @, two default categories appear: **Metrics** and **Objects**. Selecting one narrows your query to that descriptor.



Pre-Created Prompt Filters

In the Prompt field, typing a forward slash presents a list of useful pre-created prompt filters.

Copyright BE Networks 2026



You use these to preface your prompt and relegate it to specific behaviors.

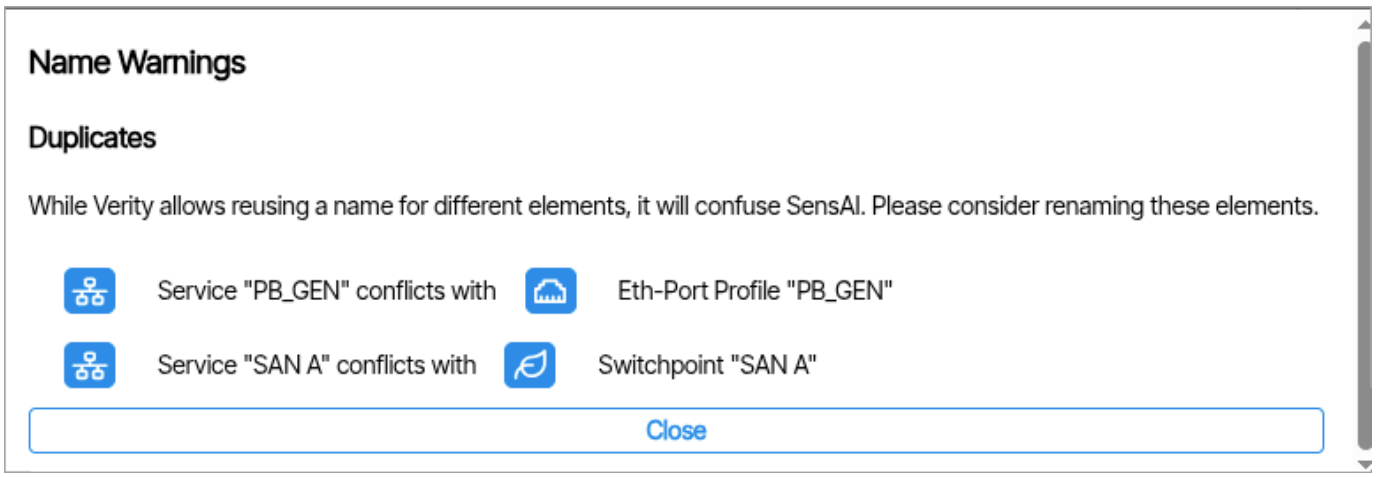


NAMING COLLISION WARNINGS

Provisioning objects with duplicate names (using the same name more than once), name collisions (names that are too similar), or common words (such as naming objects with generic terms like "a") will trigger the name collision notification icon to display yellow.



Clicking the icon opens a notification window that describes the naming issue.



Copyright BE Networks 2026

4.6.10 MCP

Client Connections

> **Connectivity Requirements**
>
> Our server must be reachable from your AI client over the network. Some agentic tools connect from the provider's infrastructure rather than your local machine, which means servers behind a corporate firewall, VPN, or private network may not be reachable even if they appear accessible from your desktop. If you're unable to connect, contact your IT or network team.

This configuration connects your MCP-compatible client to the Verity platform. The example below requires three values: Verity's server address, your username, and your password. Once configured, your client connects to the platform; the server authenticates the request and enforces access controls. The JSON structure may vary by client — this example shows the required fields for Verity.

```
{
  "mcpServers": {
    "Fabric-Manager": {
      "serverUrl": "https://{hostname}/mcp",
      "headers": {"Username": "{user}", "Password": "{pass}"}
    }
  }
}
```

MCP Terminology

Different AI tools label this configuration differently — but the setup is identical regardless. Claude Desktop and Cursor call it an MCP Server. VS Code refers to it under an MCP block. Some tools use the terms tool server, agent connector, or simply integration. If you see any of these terms in your AI software's settings, this is the configuration they're asking for.

Copyright BE Networks 2026

Telemetry

4.6.11 Telemetry Exports

Verity supports both external alert exports as well as sending time series data to external data stores. The two methods include:

- Prometheus Scraping
- AlertManager

Prometheus Scraping

The customer adds this job to **their own** Prometheus config to pull selected time-series from the Verity Prometheus instance.

```
# Customer's prometheus.yml
scrape_configs:
  - job_name: 'federation'
    honor_labels: true
    metrics_path: '/federate'
    scrape_interval: 30s
    scrape_timeout: 25s
    params:
      'match[]':
        - '{job="telegraf"}'                        # all telegraf-collected switch metrics
        - '{__name__=~"ALERTS|ALERTS_FOR_STATE"}'   # alert-related metrics
        - '{__name__="ifHCInOctets"}'               # cherry-pick a specific metric
    static_configs:
      - targets:
          - '<vendor-prometheus-host>:9090'
    # if Verity Prometheus is behind TLS or basic auth:
    # scheme: https
    # basic_auth:
    #   username: "federation_user"
    #   password: "s3cret"
    # tls_config:
    #   ca_file: "/etc/prometheus/certs/cert-ca.pem"
    #   insecure_skip_verify: false
```

KEY POINTS

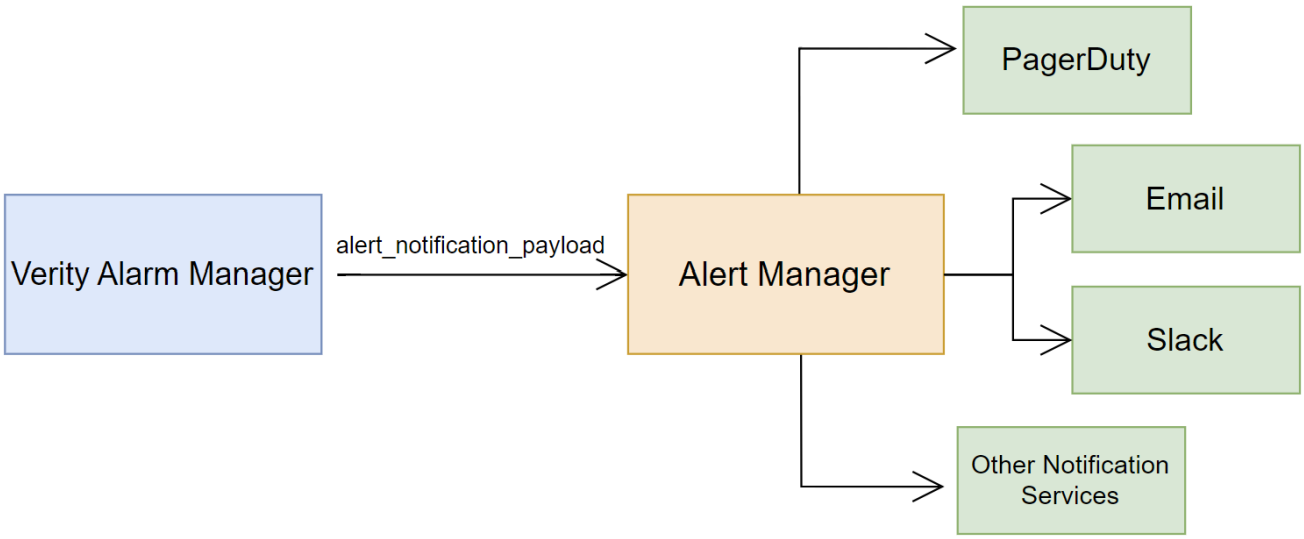
- **`honor_labels: true`** — preserves original `job`, `instance`, etc. labels from Verity instead of overwriting them with the federation job's labels.
- **`match[]`** — scope to only the series needed. Pulling everything (`{__name__=~".+"}`) is expensive and defeats the purpose of federation.
- **`scrape_interval`** — 30s is a reasonable default; should not exceed Verity's own interval (currently 15s for telegraf, 5s for sflow-rt).
- **Firewall** — port `9090` on the Verity Prometheus host must be reachable from the customer's Prometheus. A reverse proxy or VPN tunnel may be needed if the instance is not externally exposed.

Alert Manager

Alert Manager is a configurable tool that allows you to forward Verity *alarm* data to third-party messaging services such as email, PagerDuty, Microsoft Teams, and others.

The following diagram represents the relationship between Alarm Manager, Alert Manager, and an arbitrary collection of connected messaging services.

Copyright BE Networks 2026



CONFIGURING ALERT MANAGER EXPORTS

The process for editing AlertManager is:

1. Edit `/be_monitoring/alertmanager/user-config.yml` (email/receiver settings)
2. Run `/be_monitoring/utilities/render_alertmanager_config.sh` to apply

Example: Email notifications via SMTP

```
#@data/values
---
smtp:
  smarthost: "smtp.company.com:587"
  from_address: "alerts@company.com"
  auth_username: "alerts@company.com"
  auth_password: "s3cretP@ss"
  require_tls: true

notification_email:
  to: "noc-team@company.com"
  send_resolved: true
```

Example: Adding a Slack integration

```
notification_receiver_integrations:
  webhook_configs:
    - url: "https://my-ticketing-system.com/api/alert"
      send_resolved: true
```

Example: Adding a fully custom receiver + route

```
additional_receivers:
  - name: critical-pagerduty
    pagerduty_configs:
      - service_key: "abc123"
        send_resolved: true

additional_routes:
  - matchers:
      - severity="critical"
    receiver: critical-pagerduty
    continue: false
```

Applying changes

```
run /be_monitoring/utilities/render_alertmanager_config.sh
```


Copyright BE Networks 2026

This:

1. Merges vendor config + user values + overlay via `ytt`
2. Writes the final alertmanager.yml
3. Reloads Alertmanager via `POST /-/reload`
4. Persists user configs to `/be_install/archive/` for upgrade survival

SCENARIO 1 — ALERTMANAGER WEBHOOK RECEIVER BEHIND TLS

Edit `user-config.yml` (the ytt data values file). Two approaches depending on required control level.

Option A — `notification_receiver_integrations` (simplest)

Leverages built-in overlay logic. Auto-creates a receiver **and** a matching route for `job="alarm_manager_notification"` alerts, generating a receiver named `notification-webhook-0` with `continue: true`.

```
# user-config.yml
notification_receiver_integrations:
  webhook_configs:
    - url: "https://customer-noc.example.com/alerts/webhook"
      send_resolved: true
      http_config:
        tls_config:
          ca_file: "/etc/alertmanager/certs/ca.pem"
          cert_file: "/etc/alertmanager/certs/client.pem"
          key_file: "/etc/alertmanager/certs/client-key.pem"
          insecure_skip_verify: false  # set true for self-signed certs (not recommended in prod)
```

Option B — `additional_receivers` + `additional_routes` (full control)

Use when custom matchers are needed or alerts beyond `alarm_manager_notification` must be routed.

```
# user-config.yml
additional_receivers:
  - name: customer-tls-webhook
    webhook_configs:
      - url: "https://customer-noc.example.com/alerts/webhook"
        send_resolved: true
        http_config:
          tls_config:
            ca_file: "/etc/alertmanager/certs/ca.pem"
            cert_file: "/etc/alertmanager/certs/client.pem"
            key_file: "/etc/alertmanager/certs/client-key.pem"
            insecure_skip_verify: false
          # basic_auth:          # optional if webhook requires it
          #   username: "vendor"
          #   password: "s3cret"

additional_routes:
  - matchers:
      - severity=~"critical|warning"
    receiver: customer-tls-webhook
    continue: true
```

Mount TLS certs into the AlertManager container

Add a volume bind in the Docker Compose override so cert files are accessible at the paths referenced above.

```
# docker-compose override
alertmanager:
  volumes:
    - /path/to/customer/certs:/etc/alertmanager/certs:ro
```

Re-render and reload after editing

1. Render the final `alertmanager.yml` via ytt:

```
./utilities/render_alertmanager_config.sh
```

2. Reload AlertManager (or restart the container):

```
curl -X POST http://localhost:9093/-/reload
```


Copyright BE Networks 2026

Enabling TLS Encryption

Reference docs:

- http_config
- tls_config
- Federation
- TLS encryption guide

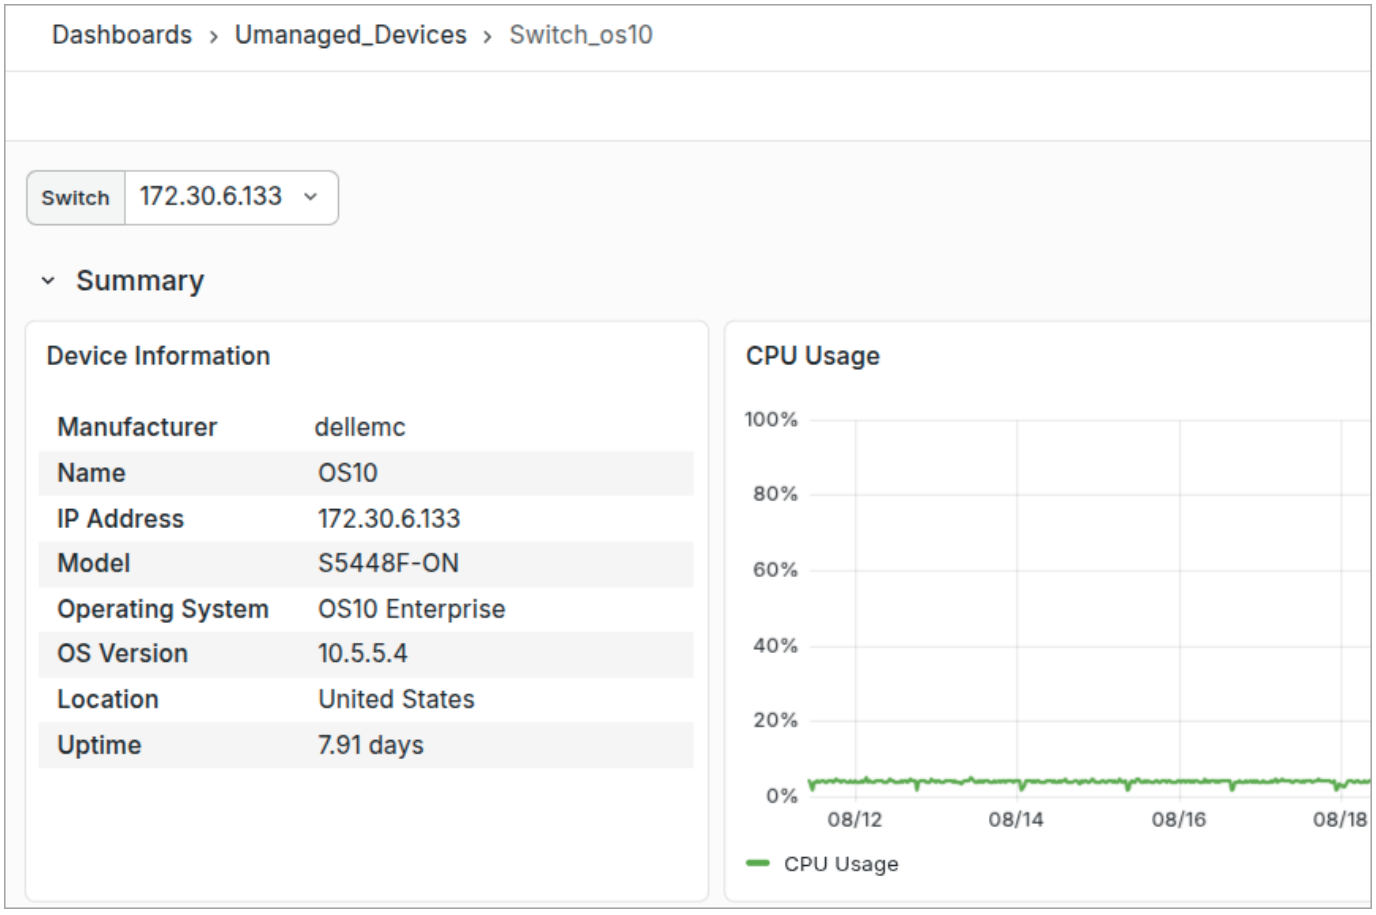
Copyright BE Networks 2026

4.6.12 Dell OS10 Monitoring

Verity supports monitoring of Dell OS10 devices via SNMP. Using the standard SNMP methods, Verity will regularly poll added devices and display the gathered statistics in device dashboards.

Configuring OS10 Monitoring

1. Configure OS10 Switches with a "snmp_importer" read-only community string using the switch device terminal.
2. Monitoring accepts a plain-text file named "os10.txt" containing a single IP address per line (Switch IPs only). Upload it in the Verity GUI via **Admin → Config File** under **Vendor Config**. After a poll, the system updates using those IPs and fetches SNMP metrics from the switches.
3. Navigate to **Observe → Unmanaged Devices → Switch_OS10** and **Port_OS10** to verify the IPs in the dropdown and view metrics for each device.



Copyright BE Networks 2026

4.6.13 Dell OME Integration

Verity can connect to a Dell OME server to gather telemetry and display server information in the automated dashboards. Verity routinely polls the OME API with Read-Only access and places the inventory and telemetry into the Verity datastore.

> **Note**
>
> Servers added to Dell OME MUST have Dell OME Enterprise Advanced licenses in order to provide the necessary telemetry to Verity.

Configure Dell OME Server Integration

1. Go to Admin > Dell OME.
2. Add a new server and provide the necessary credentials to reach the server.
3. Dell OME dashboard will be created in a few minutes.
4. Navigate to the the dashboard and select the server Host OS name from the dropdown list at the top.

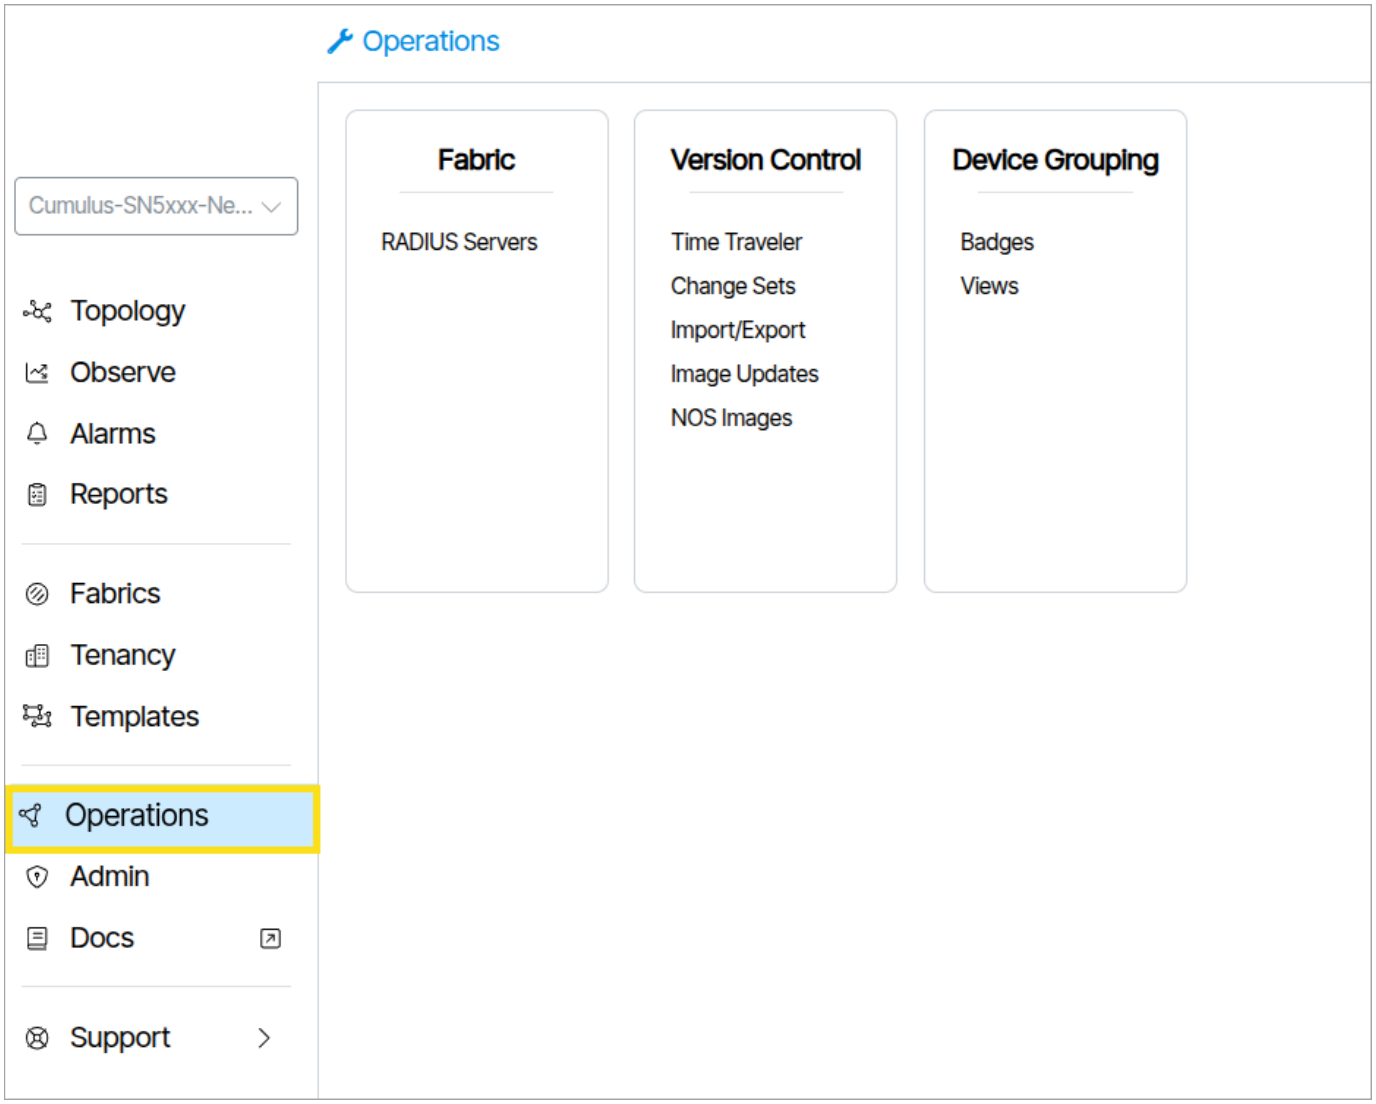
Copyright BE Networks 2026

4.7 Operations

Operations

4.7.1 Operations

Operations provides tools that enable administrators to manage existing network infrastructure.



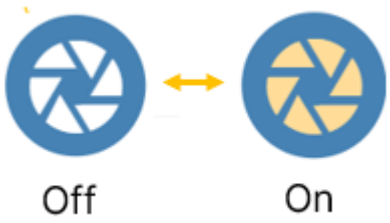
Copyright BE Networks 2026

4.7.2 Network Operations

There are a number of common operations that operators may perform on their network.

Network Read-Only Mode

1. Click **Topology** and navigate to the **Network View**, click the **Read-Only** icon to "on" ( )



> **Note**
>
> The read-only mode can be applied to the world and per switch. All levels are honored simultaneously and are displayed on each switch.
>
> World wide read only is automatically set during a CORE upgrade of the vNETC. It is recommended, but not required, to be set for firmware and other major system updates to allow the user to see what unexpected changes may be caused by those actions.

Switch Replacement

The new device should be in ONIE mode.

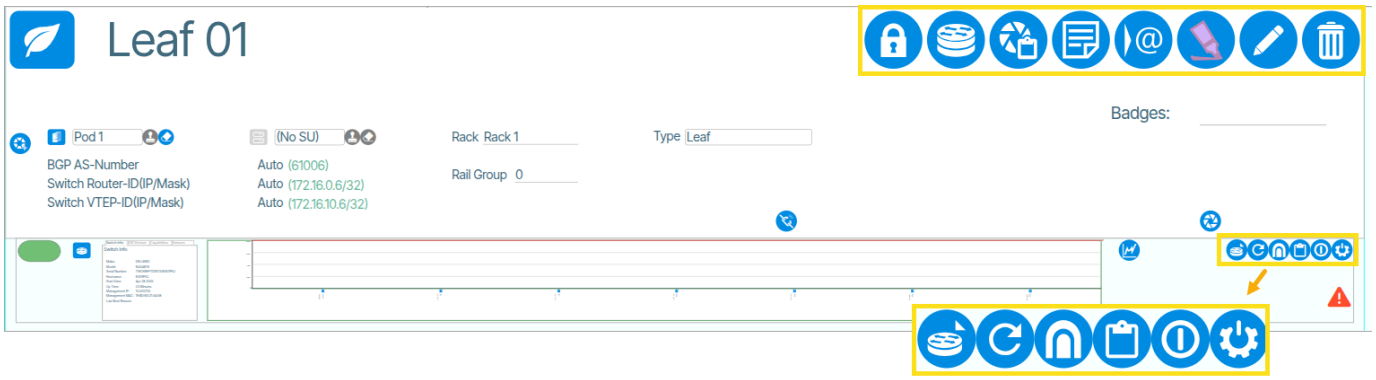
1. Remove the old device from service. ( )
2. Delete the old device. ( )
3. In the option box, make sure "Preserve Switchpoint" is selected. ( )
4. In the Preprovisioned Switchpoint, clear the device ID field ( ).
5. Go to the **Device Management** settings (Fabrics → Device → Device Management).
6. Update the ZTP Identification field.
7. Re-enable the Device Management.

Refer to the ZTP sections for the ZTP process for the switch role (ie Leaf, Spine or Management). The new switch shhould go through ONIE and ZTP process and start communicating with the Device Controller.

4.7.3 Device Operations

Inspecting devices and performing operations on them is accomplished in the Switchpoint object in the Topology view.

Copyright BE Networks 2026



Using Device Operations

Zoom to a device by double clicking on it in the topology map. ()

HOW TO REBOOT A SWITCH

From **Topology**, zoom into the network operations section (top lower right) of the device and click the reboot button.



DEVICE LOCKS



Locking a device restricts users without **[NW] Network** privileges from editing it.()

TO LOCK A DEVICE:

1. Click **Topology**
2. Zoom in to the top of the device. 
3. Click the lock icon ().

Copyright BE Networks 2026

DEVICE OPERATIONAL TOOLS

Copyright BE Networks 2026

Icon	Name	Function
	Show the current switch configuration	Show current switch configuration
	MAC Address Search	Search for a MAC address on this device
	System Log	View the system log
	Highlighter	Highlight the device on the topology map
	Edit	Edit the device settings
	Lock	Locks the device
	Delete	Delete object
	Manage Switch Pairs	Add Switch Pairs
	Opens Ping/Traceroute Dialog	Open Ping/Traceroute Dialog

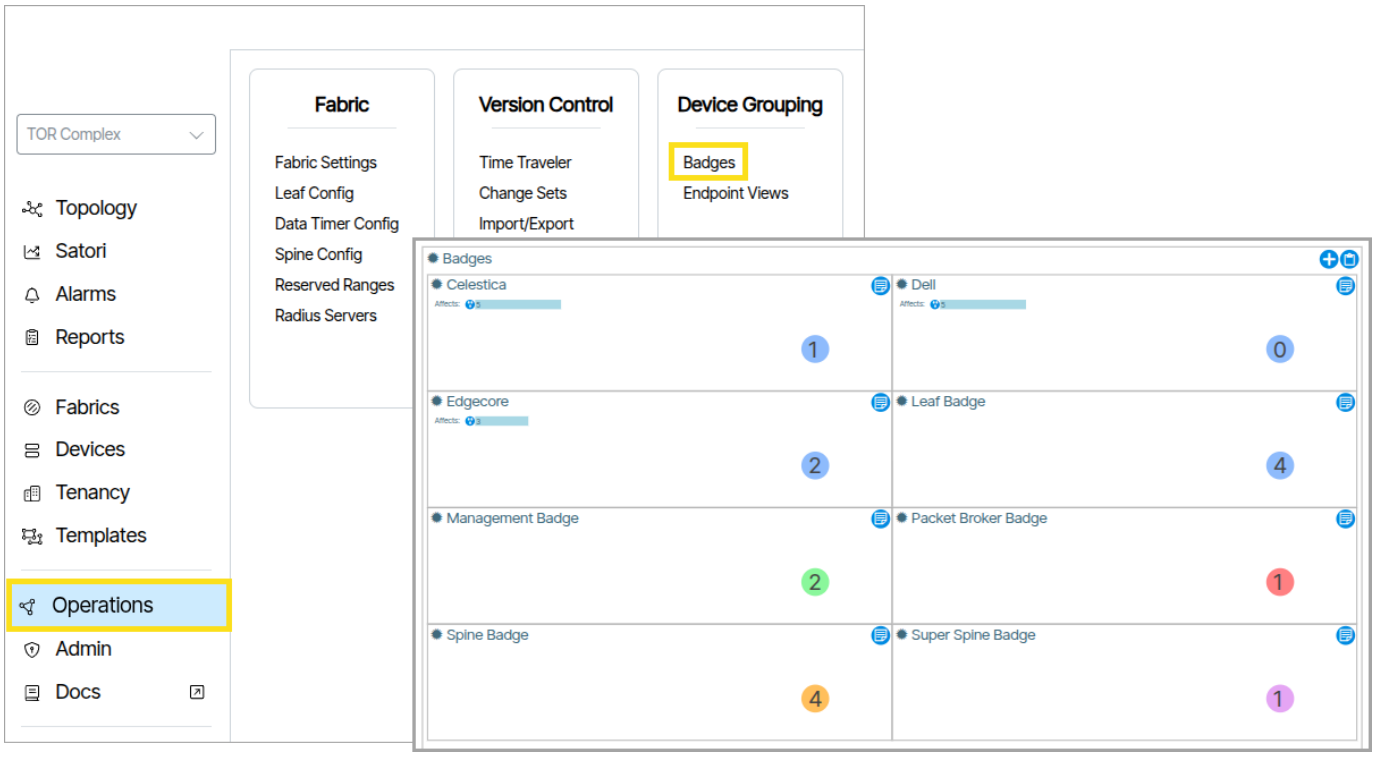
Copyright BE Networks 2026

Icon	Name	Function
	Capture Snapshot	Capture a snapshot of this device for troubleshooting - press once to start capturing and then again when capture is complete (icon turns dark green)
	Rescan	Trigger a full rescan by the ACS of this device.
	Open Tunnel	Open a remote access tunnel
	Report	Open report
	Reboot Device	Reboots Device
	Out of Service	Mark device out of service.

Badges

Badges are used to group devices visually. Badges are assigned a name, a color and a number. To create a new badge or view all badges go to **Operations/Badges**.

Copyright BE Networks 2026



CREATE AND ASSIGN BADGE TO DEVICE

1. Click **Topology**.
2. Enable focus of a selected device.
3. Click the **Edit** icon.
4. Click set badges icon ( ).
5. Enable the check-box next to the phrase **Create New Badge**.
6. Set a **Name**, **Color** and **Num** for the badge.
7. Click **Yes** to create the badge.

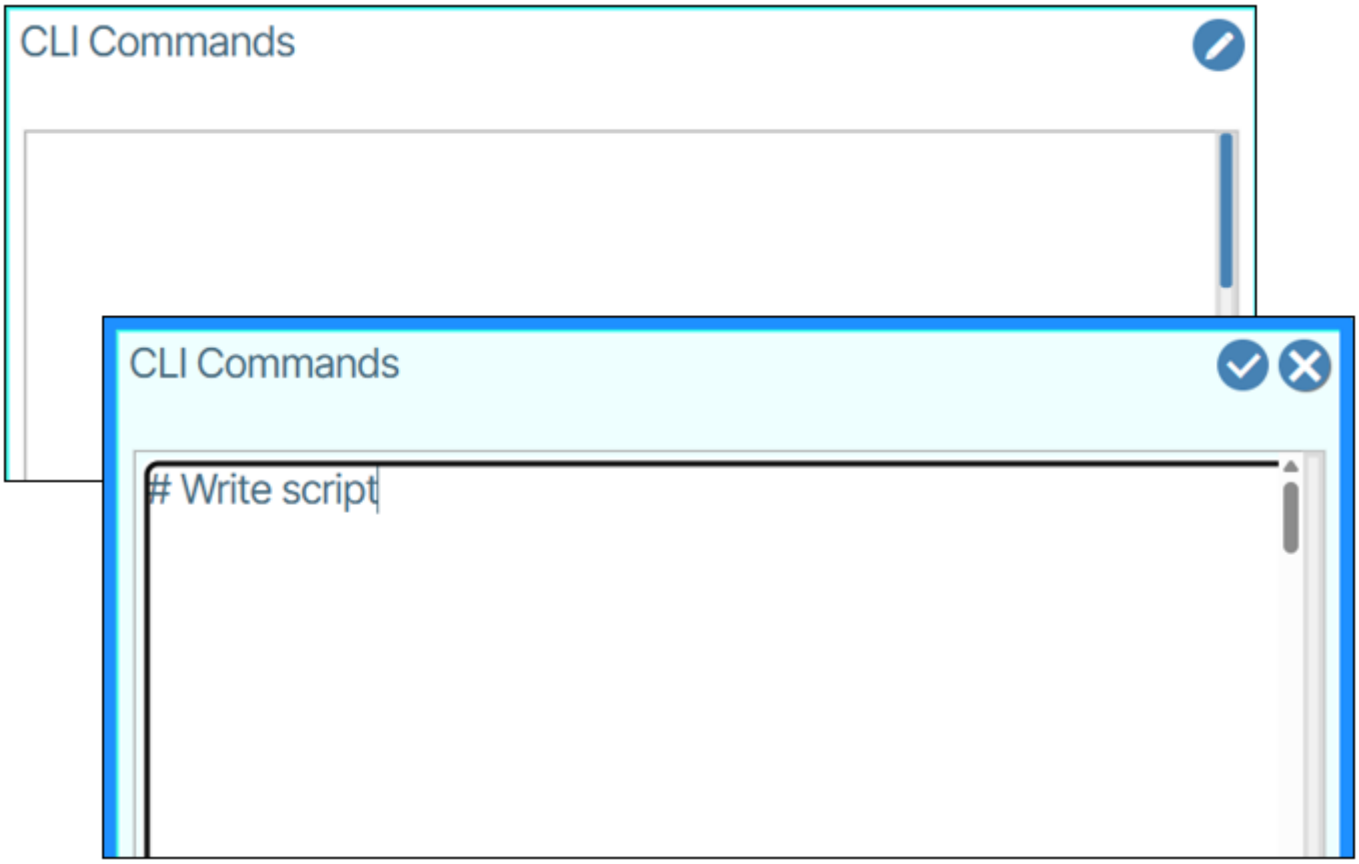
Copyright BE Networks 2026

4.7.4 Device CLI

There are three CLI command interfaces, each designed for inputting commands into the switch. Each interface impacts the device in a distinct manner.

- **Global/Switch Wide**: Commands are applied once at the global configuration level of the switch ( ).
- **Service/Application**: Commands are applied against the corresponding vlan at the vlan interface level ( ).
- **Port Settings**: Commands are applied at the interface level ( ).

To write commands, click the edit ( ) button and type your commands. Clicking the save button ( ) transmits the commands.



Warning

When Submitting CLI Commands:

- Commands are not checked for syntax and are submitted exactly as entered.
- Commands are reapplied whenever the text is changed or when the controller reconnects to the switch.
- The commands are not compared to the running configuration, so they will appear in pending changes until they are first applied by the controller.
- Changing of any parameters controlled by Verity may be immediately reverted by the controller.

Target Specific CLI Information:

- **Dell/Broadcom SONiC**: Uses sonic-cli and starts at the global configuration mode.

Copyright BE Networks 2026

- **Edgecore SONiC**: Commands are executed at the Linux shell level. Please note that only the Global CLI text is supported.
- **Cisco**: Commands are at the global configuration level.

Copyright BE Networks 2026

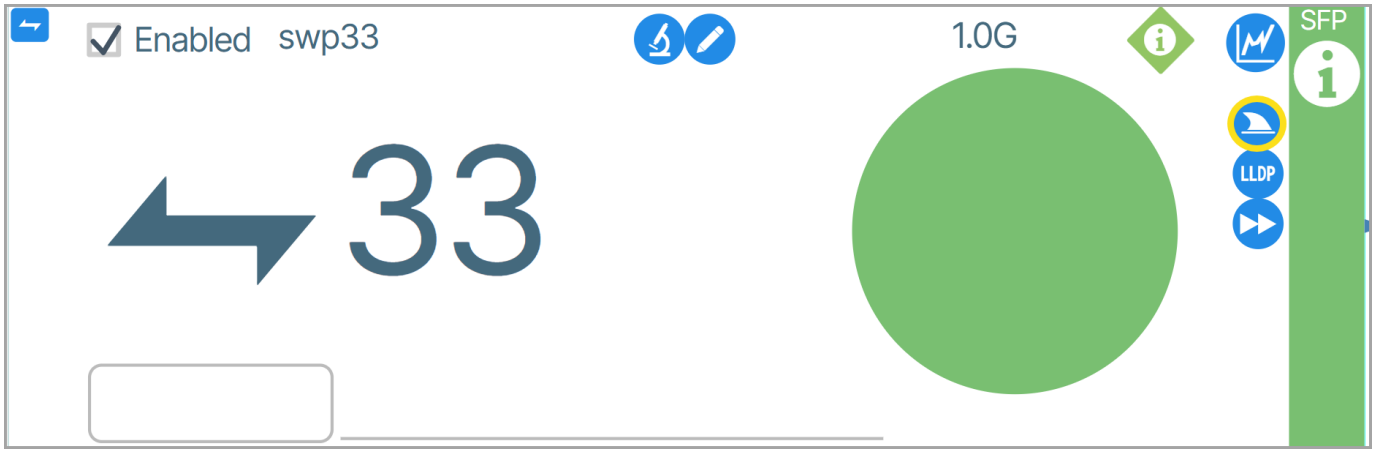
4.7.5 Packet Capture

PACKET CAPTURE

Packet Capture is only available on select hardware

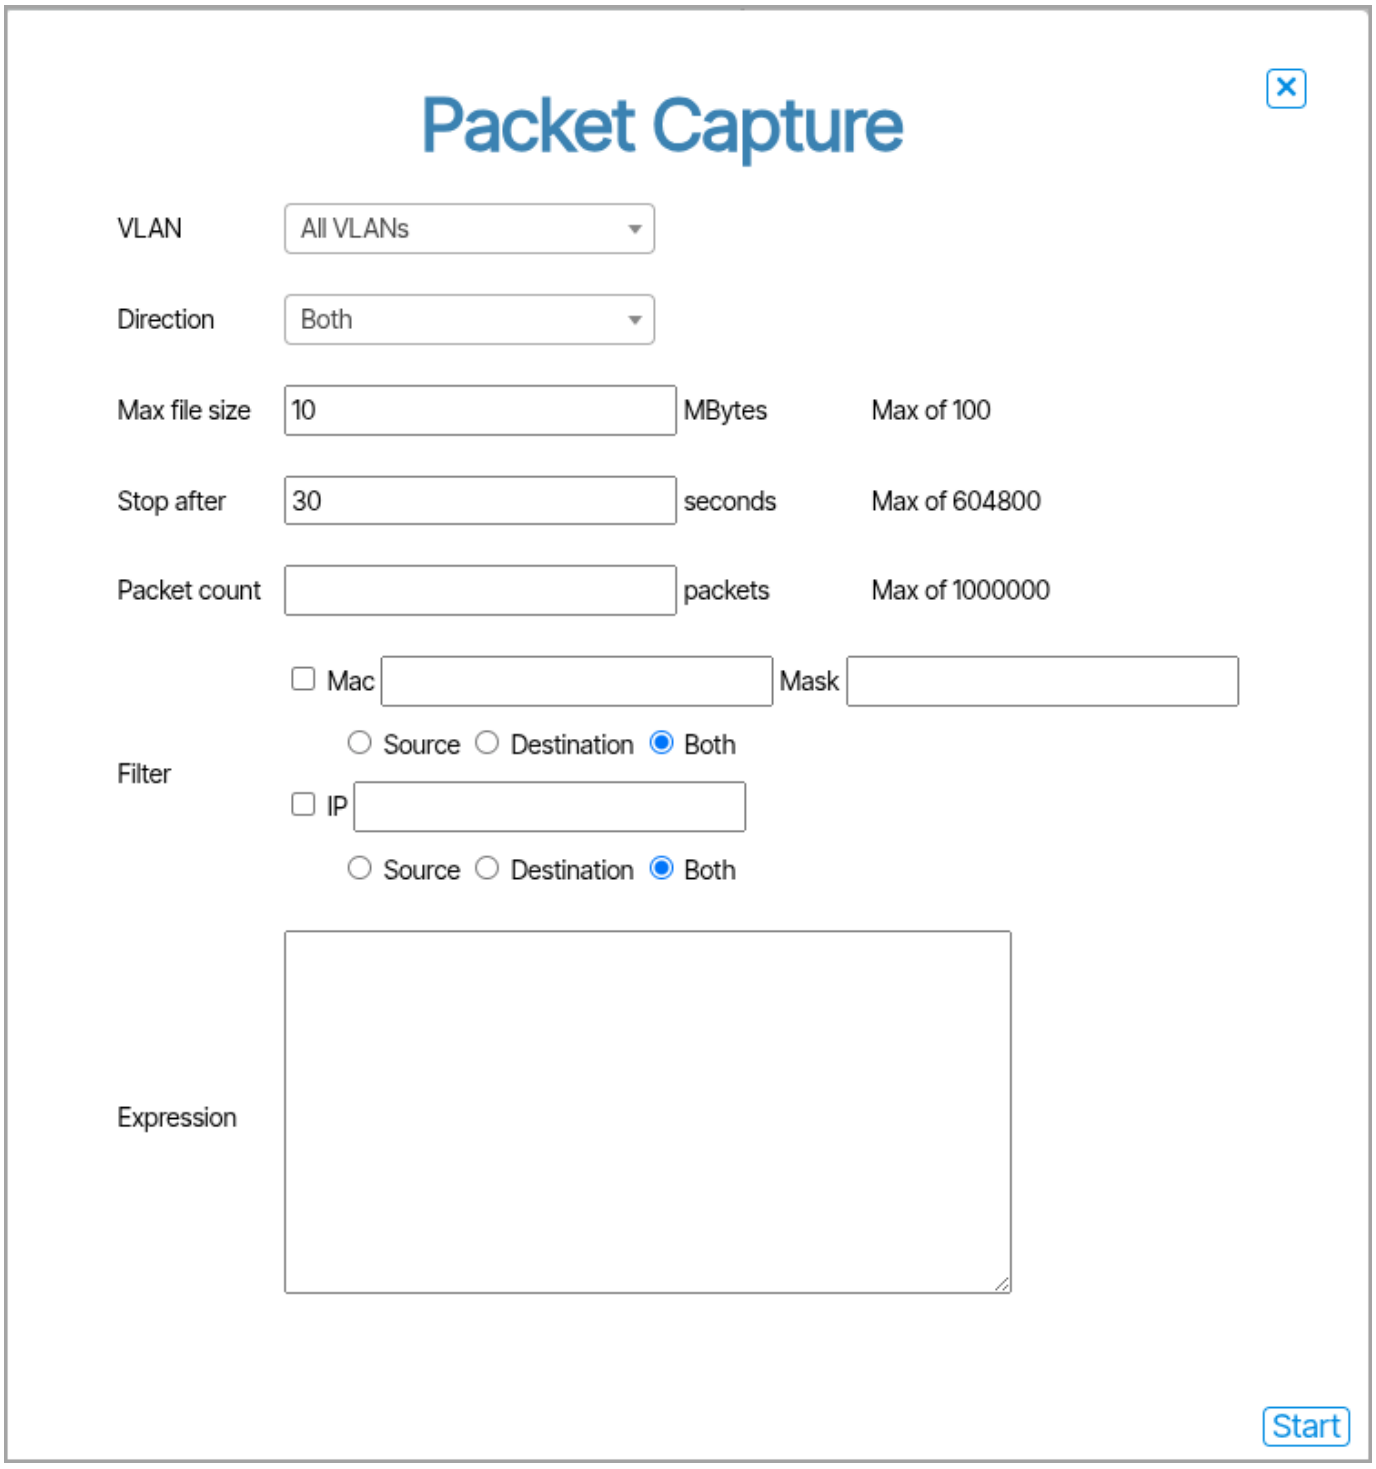
This feature lets you capture network packet information and save the data to a file.

To start a packet capture, click the shark fin icon ( ).



The **Packet Capture** window appears. Enter the parameters for the data you want to capture and click **Start.**

Copyright BE Networks 2026



When data capturing is complete, a *pcap* file containing the information will be downloaded to your computer.

Copyright BE Networks 2026

4.7.6 Troubleshooting

Topology

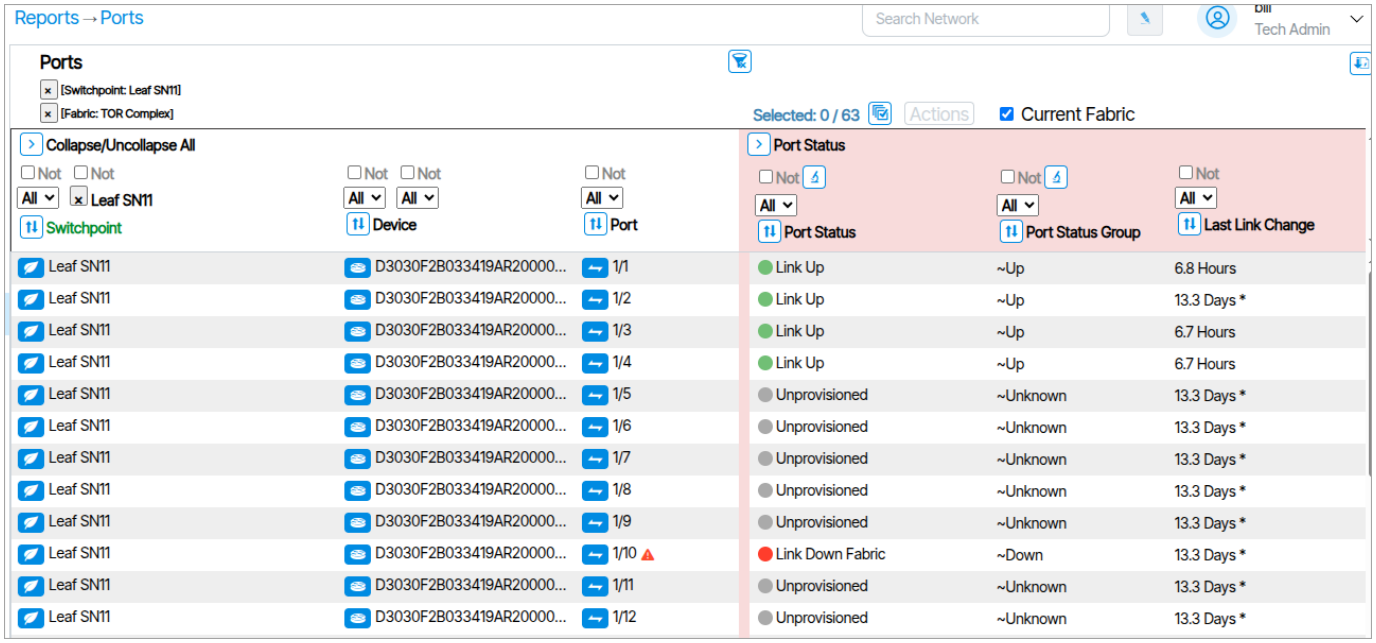
Verity displays a live map of the devices in the network and their connections.

TOPOLOGY DETECTION AND BEHAVIOR

1. Autodiscovery
2. Connection moves
3. Remove connections

Reports

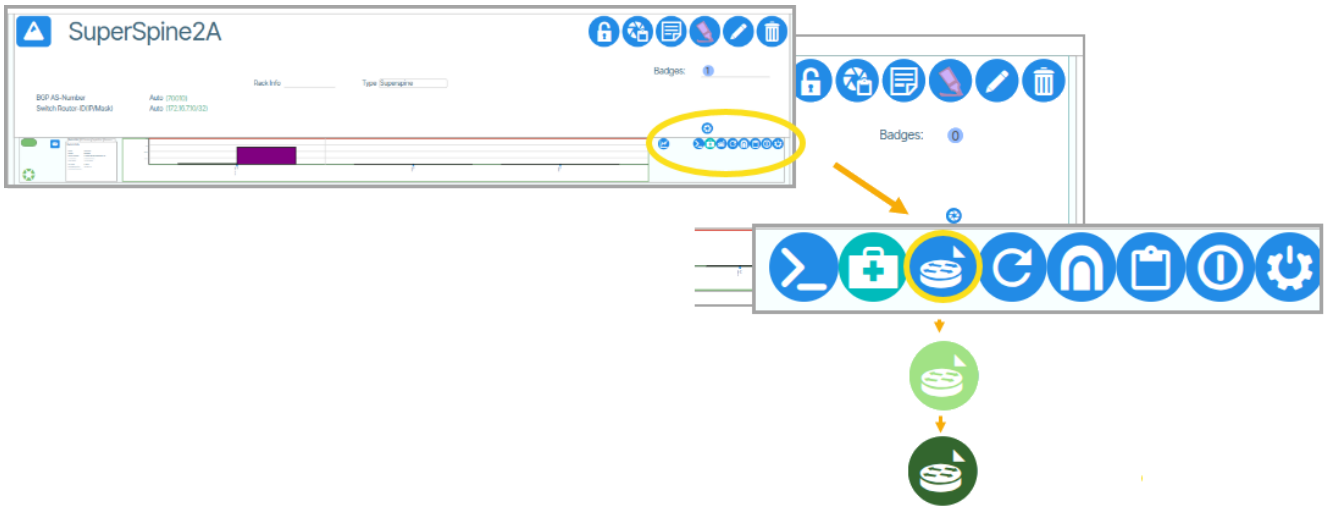
Verity provides a comprehensive list of reports to identify current conditions in the network. Reports are viewed by navigating to the Reports window  or clicking the Reports icon of a device object  .



HOW TO DOWNLOAD A DEVICE SNAPSHOT

To download switch device information, locate the device in the topology view. In the device's top-right corner, click **Capture Download Snapshot**. The button turns light green while the snapshot is being prepared, then dark green when it's ready. Click it again to download.

Copyright BE Networks 2026

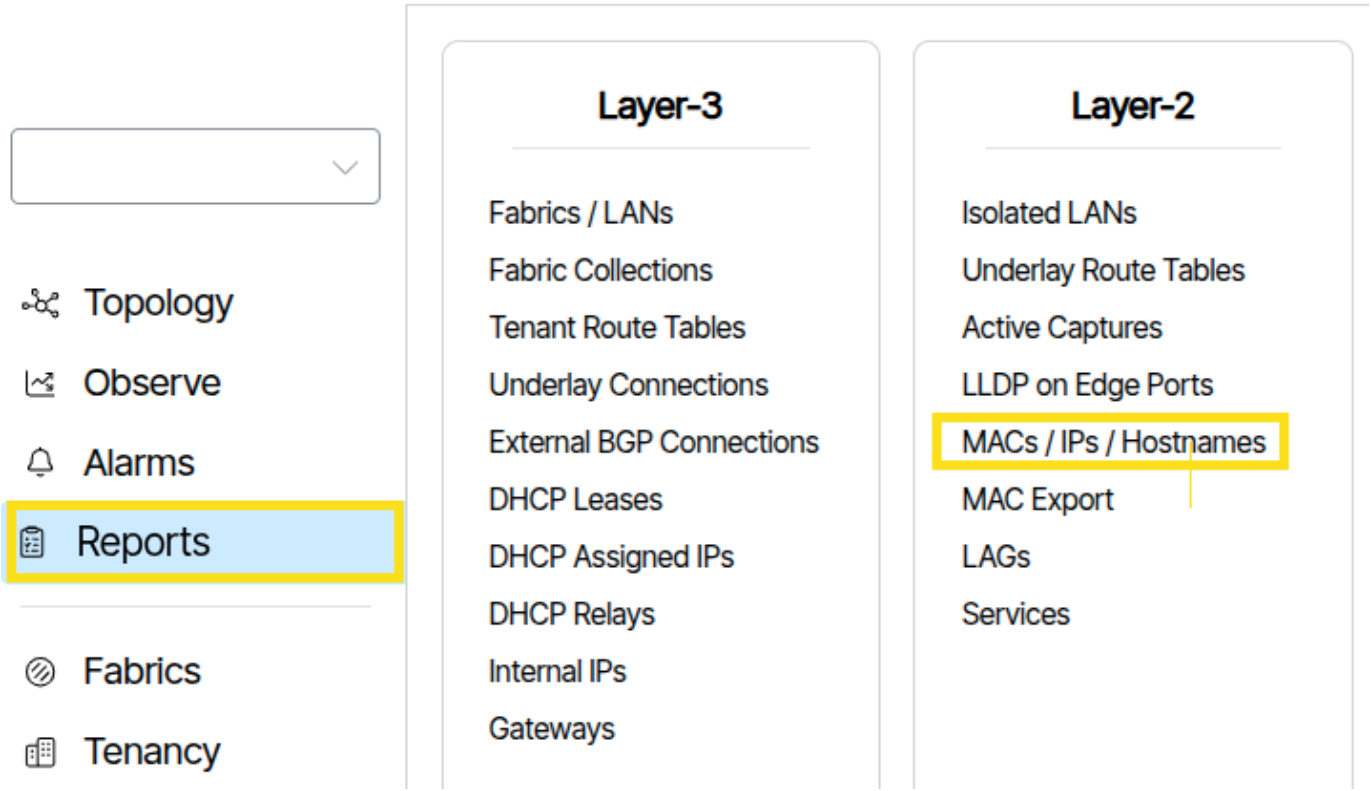


HOW TO DOWNLOAD A VERITY SNAPSHOT

In addition to individual device snapshots, Verity allows for download snapshots of VNC, Observability VM, AI Ops VM, and GUI software. These are used for broader troubleshooting of the host software. Instructions are available here.

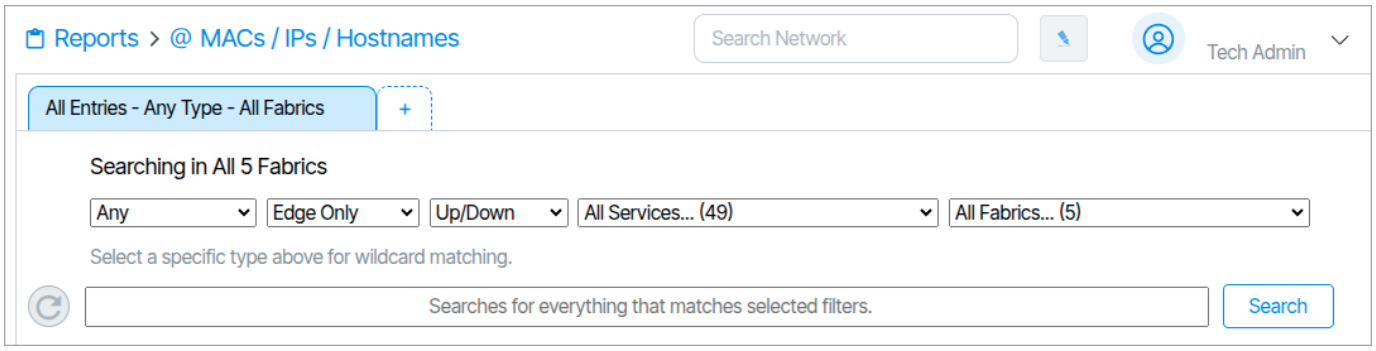
MACs / IPs / Hostnames

Verity provides the ability to display MAC addresses, IP's and hostnames across the entire system. To use this feature navigate to **Reports/ MACs / IPs / Hostnames**.



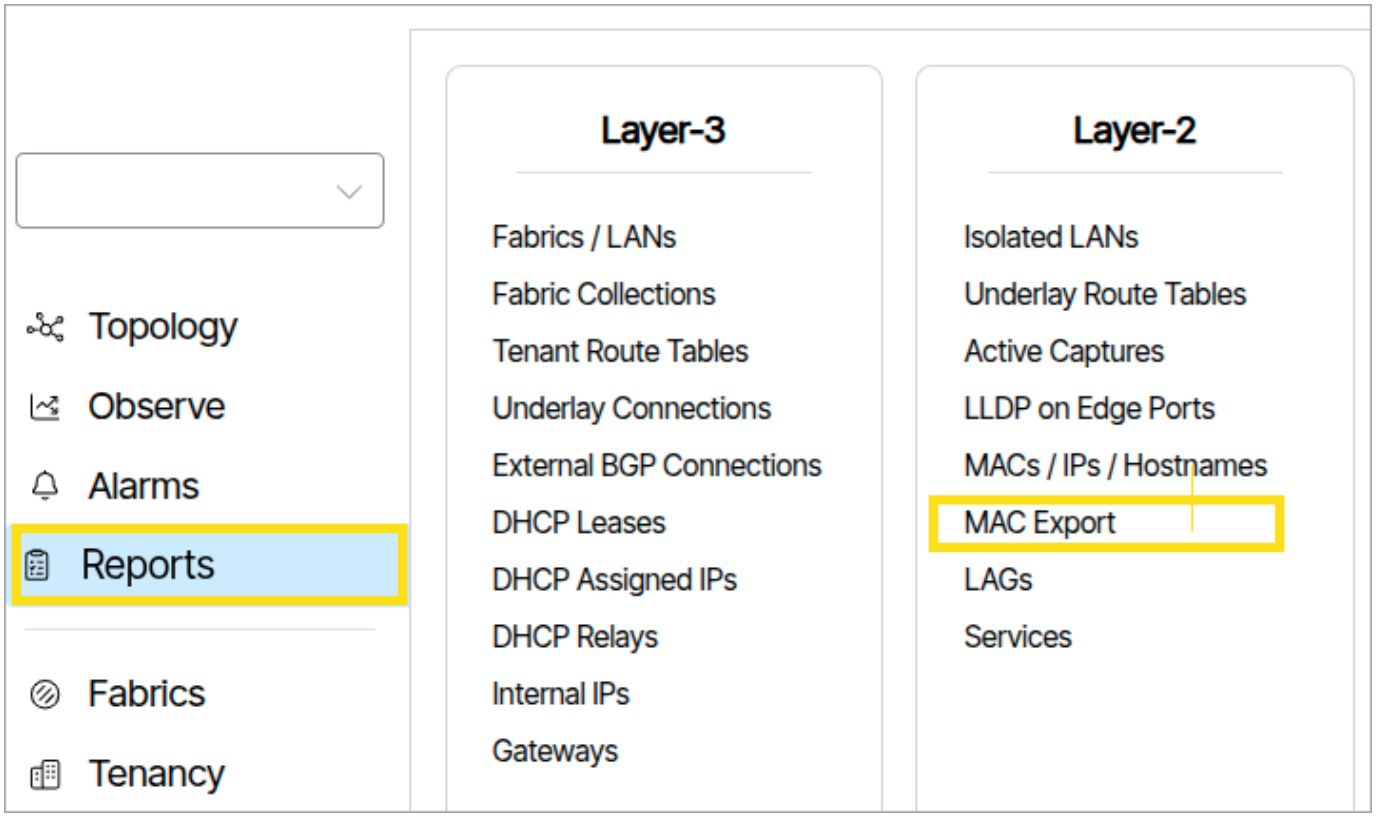
The search field defaults to **MAC Address** and can be changed to search for IP address or hostname.

Copyright BE Networks 2026

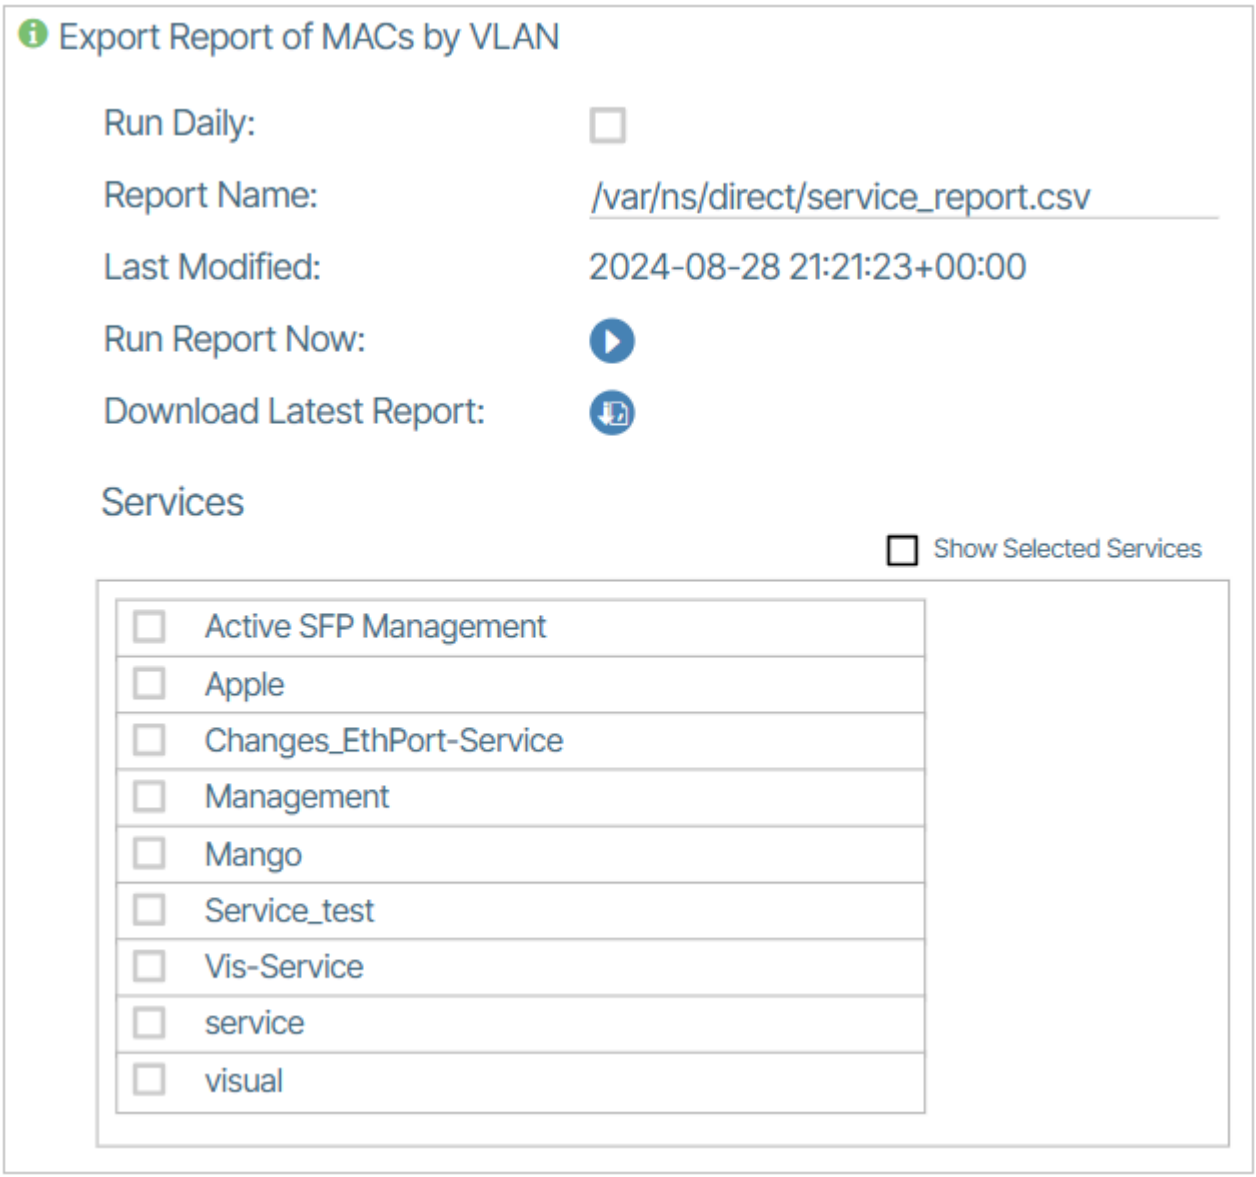


Mac Exports

Mac Exports provides a summary MAC report for all devices connected to the network's user facing ethernet ports. This tool can be automated to generate a daily report or used on a case-by-case basis as needed. To manually run a report, enable the checkbox for all services you want to include, click **Run Report Now**, and wait for the **Download Latest Report** icon to turn blue. Once it does, click the icon to download the CSV file.



Copyright BE Networks 2026

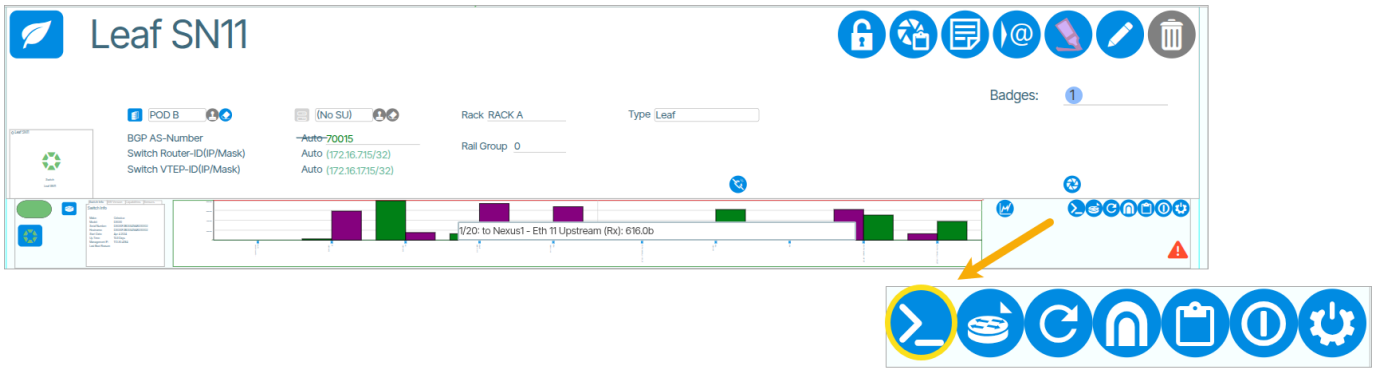


Ping and Trace Route

Verity provides tools to troubleshoot Layer 3 issues for the underlay and Tenant connections for services configure for Layer 3. The **Ping and Trace Route** feature is a diagnostic tool that sends a 'ping' to and from a specified address and details the message path.

CONTROLS

The **Ping and Trace Route** dialog box is opened by clicking **Open Ping/Traceroute Dialog** button .



Copyright BE Networks 2026

To use this feature you configure the following parameters and click the **Start** button .

SOURCE

- **Device:** This is the source device of the ping.
- **VRF:** This is where you select the VRF. The VRFs of Tenants, the Underlay (selected device), and Management (Device Controller) are all options.
- **IP:** The targeted IP address to send the ping command from. The IP address changes depending on the VRF setting.

DESTINATION

- **IP:** The targeted IP address to receive the ping command.

PING

- **Count:** The number of ping events sent before the process is complete.

TRACE ROUTE

When this option is selected, the trace route is included in the **Result** field.

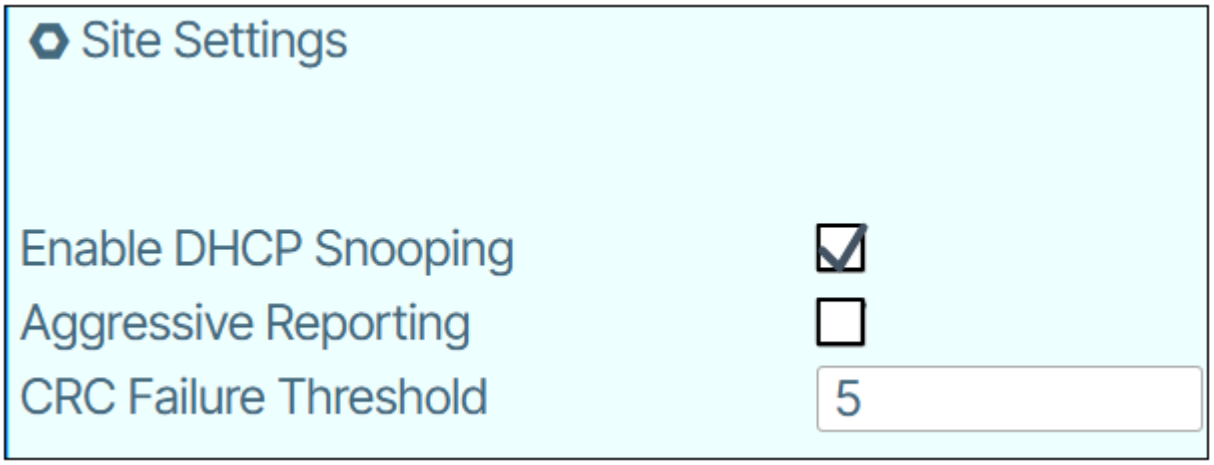
CMD STATUS

This is a progress indicator. When the feature is ready, CMD Status says **Ready**; when active, CMD Status says **In Progress**; and when the process is complete, it says **Done**.

DHCP Snooping

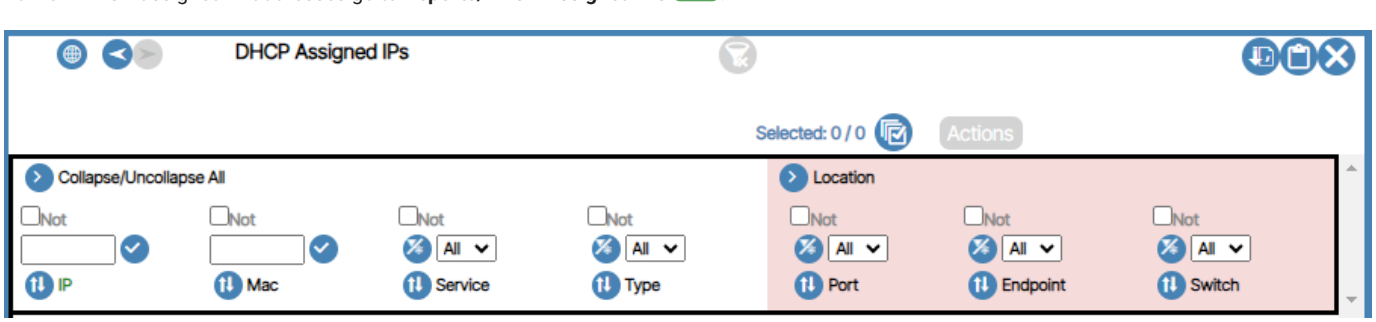
Verity provides the location of connected L3 addresses detected by DHCP snooping functions in SONiC.

To enable DHCP snooping navigate to **Admin/Site Settings** and check the box next to **Enable DHCP Snooping** .



VIEWING DHCP ASSIGNED IP'S

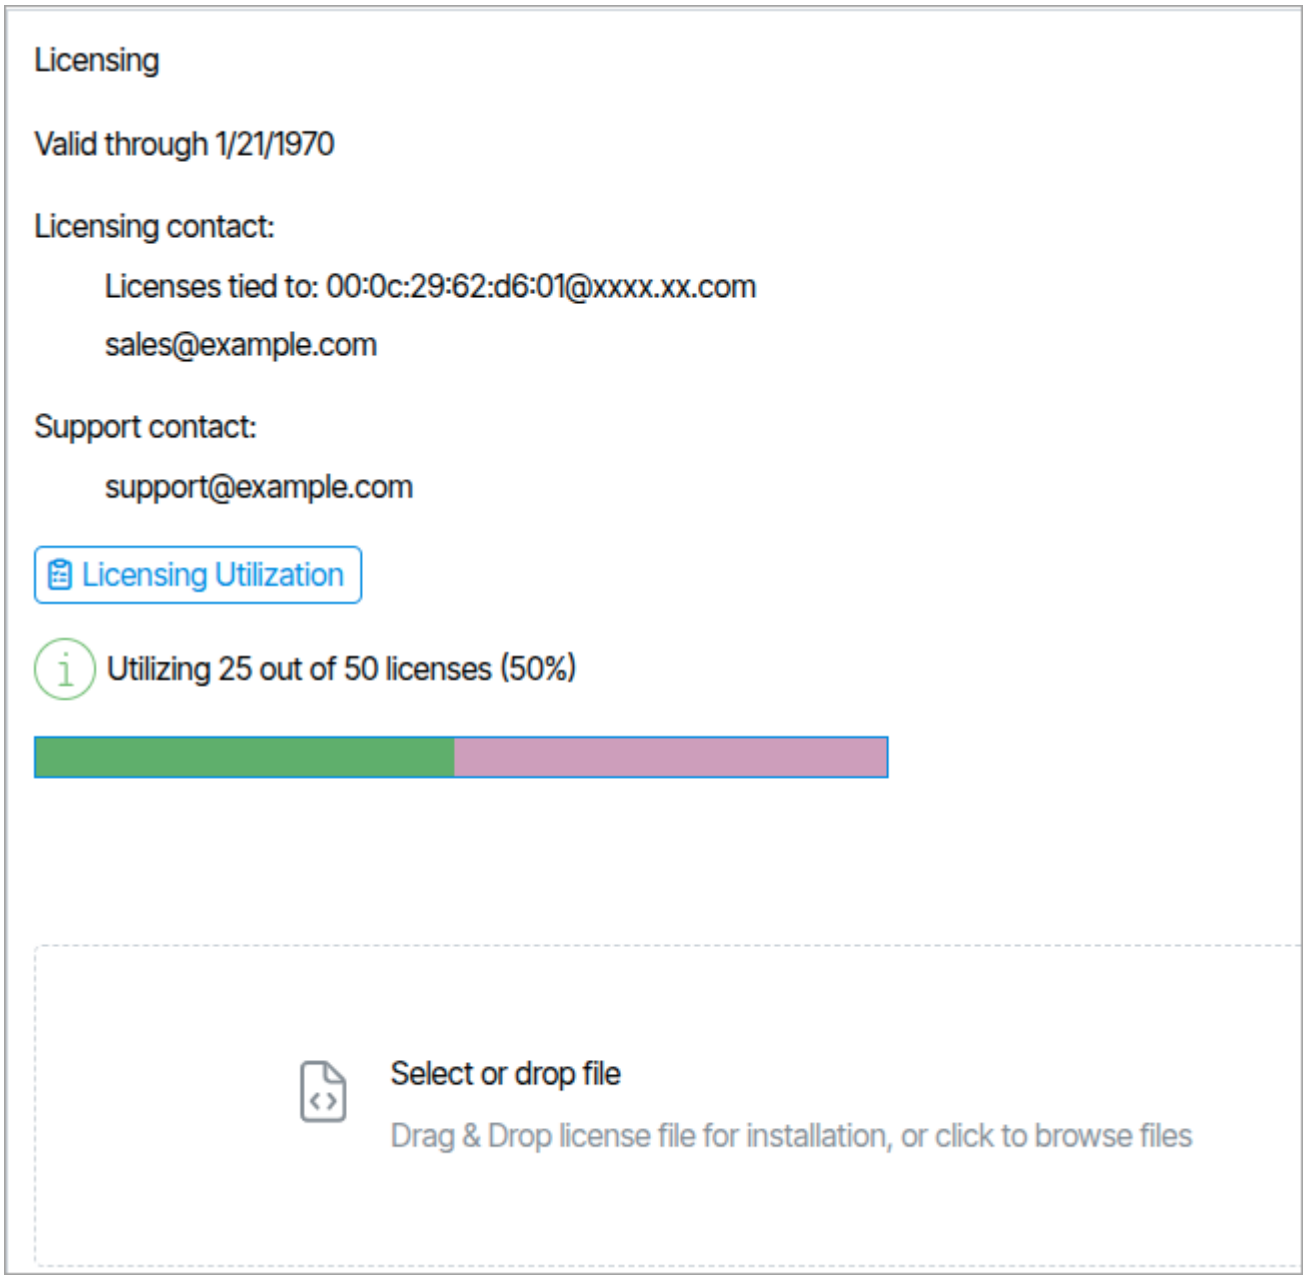
To view DHCP assigned IP addresses go to **Reports/DHCP Assigned IPs** .



Copyright BE Networks 2026

View License

To view the license go to **Admin/Licensing**.

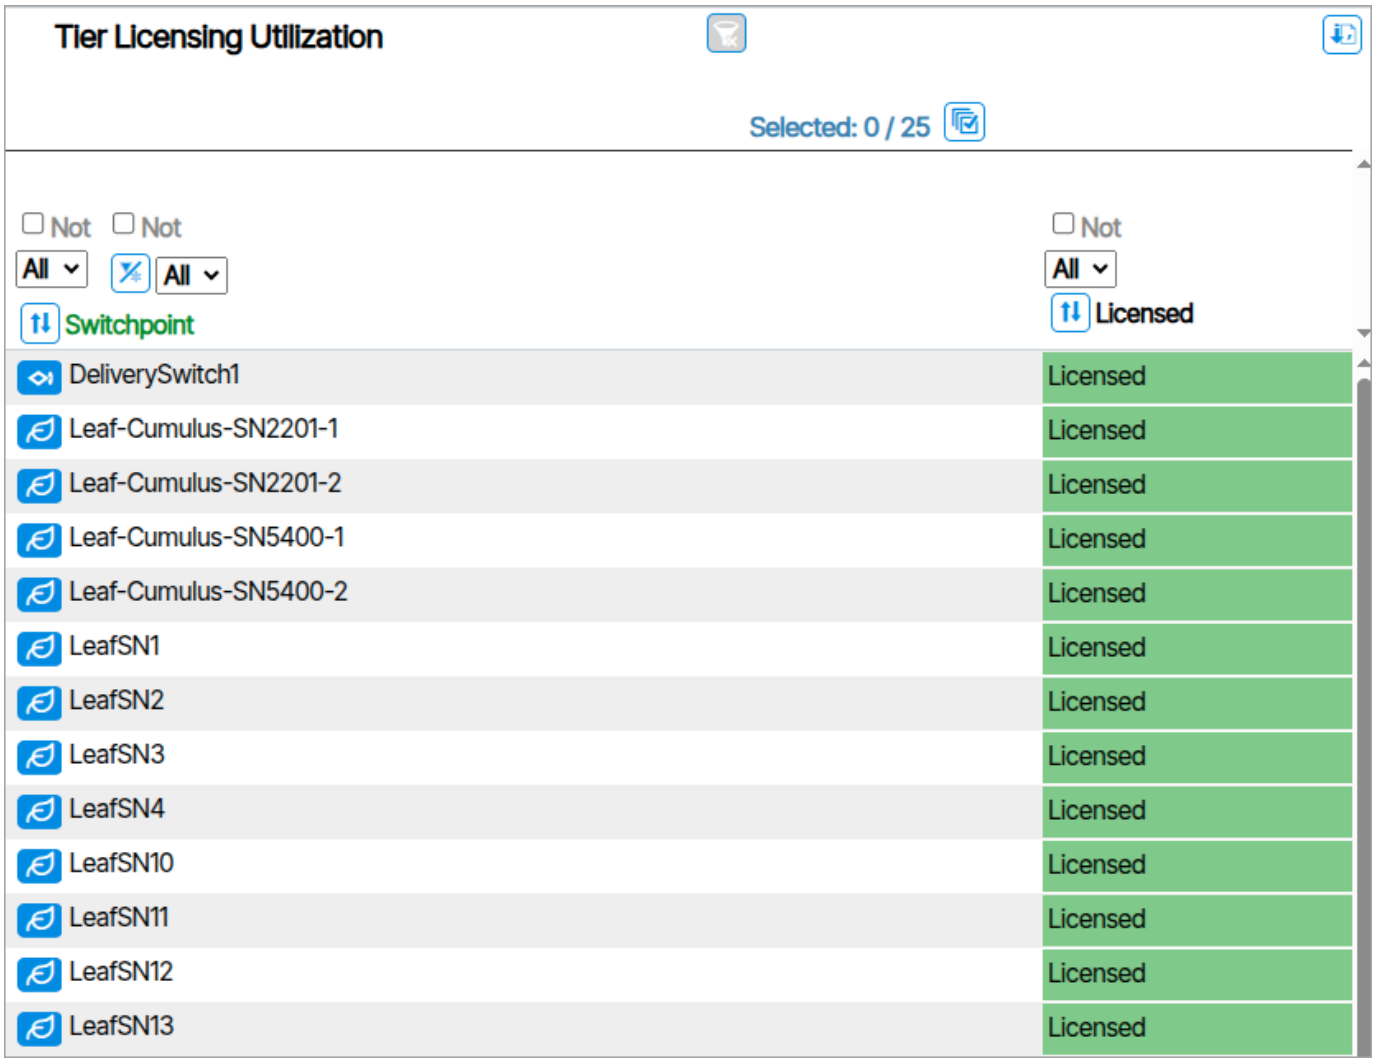


The Licensing object displays License Usage and Physical Port Usage bar graphs. While focused on the object, the user can see the date of licensing expiration, contact information for support, and reports of license.

By clicking on the report icons for License Utilization a report for thh list will be shown. For a specific report on used, preprovisioned, stranded, or spare licenses, the user can click on each section of the bar graph. This can also be done with physical ports that are licensed, fabric or spare.

An example License Utilization report is shown below:

Copyright BE Networks 2026



4.7.7 Monitoring Troubleshooting

Overview

This troubleshooting script is a comprehensive diagnostic and maintenance tool for the Verity Monitoring System. It provides a menu-driven interface to check system status, perform maintenance tasks, and generate debug information for support purposes.

Purpose

The script helps administrators:

Diagnose Issues: Check service status, connectivity, and configuration

Perform Maintenance: Adjust logging levels, restore backups, populate data

Generate Support Data: Create comprehensive debug snapshots for technical support

Monitor Health: Verify all Monitoring components are functioning correctly

Prerequisites

SYSTEM REQUIREMENTS

Privileges: Must run as root or with sudo access

Copyright BE Networks 2026

Dependencies:

- Monitoring system previously installed via setup script
- Active internet connection for API testing

Running the Script

BASIC EXECUTION

```
bash

sudo ./monitoring_admin.sh
```

You are presented with two options:

1. Setup
2. Troubleshooting

Select 2 **Troubleshooting**.

Menu Options Reference

OPTION 1: CHECK MONITORING SERVICES STATUS

```
1. Check if the Monitoring services are running
```

Purpose: Verifies all Monitoring Docker containers are running properly

What it checks:

- ETL (fabric-manager-ml-python)
- Grafana
- Prometheus
- Promtail
- Telegraf
- Neo4j-graph
- Loki
- Cadvisor
- Alertmanager
- Node-exporter
- AI-service
- Neo4j-vector

Expected Output:

```
ETL services are running.
Grafana services are running.
Prometheus services are running.
...
---------------------------------------------
Total Services Tested: 12
Currently Running: 12
Not Running: 0
---------------------------------------------
```

If Services Are Down:

```
run the command 'sudo docker compose -f /be_monitoring/docker-compose.yml up -d' and re-run this script
```


Copyright BE Networks 2026

OPTION 2: CHECK TIME SYNCHRONIZATION

```
2. Check the status of the time synchronization
```

Purpose: Verifies system time is properly synchronized

Expected Output:

```
               Local time: Wed 2024-12-18 10:30:15 UTC
           Universal time: Wed 2024-12-18 10:30:15 UTC
                 RTC time: Wed 2024-12-18 10:30:15
                Time zone: UTC (UTC, +0000)
System clock synchronized: yes
              NTP service: active
          RTC in local TZ: no
```

Key Indicators:

- `System clock synchronized: yes` - Time is properly synced
- `NTP service: active` - Time synchronization service is running

OPTION 3: REPORT FQDN SETTING

```
3. Report the FQDN setting in the Monitoring.ini file
```

Purpose: Shows the configured vNetC host address

Expected Output:

```
FQDN setting: vnc-monitoring.company.com
```

Uses: Verify Monitoring knows where to connect to vNetC

OPTION 4: REPORT MAC ADDRESS SETTING

```
4. Report the FABRIC_MANAGER_MONITORING_VM_MAC_ADDRESS setting in the Monitoring.ini file
```

Purpose: Shows the configured MAC address for VM identification

Expected Output:

```
FABRIC_MANAGER_MONITORING_VM_MAC_ADDRESS setting: aa:bb:cc:dd:ee:ff
```

Uses: Verify unique VM identification in the Monitoring system

---

OPTION 5: REPORT PACKAGE VERSION

```
5. Report the version of the Monitoring package
```

Purpose: Shows the currently installed Monitoring software version

Expected Output:

```
Version: 2.1.64.108
```

Uses: Verify version for support purposes or update planning

OPTION 6: PERFORM ANNOUNCEMENT CALL

```
6. Perform an announcement call to the FQDN setting
```

Purpose: Tests Monitoring's ability to announce itself to vNetC

Copyright BE Networks 2026

What it does: Executes the announcement process to register with vNetC

Expected Behavior: Script runs the announcement process and reports any failures.

OPTION 7: CHECK VNETC API STATUS

```
7. Check vNetC API calls are working and report the HTTP status code
```

Purpose: Tests API connectivity to the vNetC system

Expected Output (Success):

```
Verity vNetC API calls are working using the FQDN/ROOT_API_URL: vnc-monitoring.company.com
```

Expected Output (Failure):

```
Verity vNetC API calls are not working. FQDN/ROOT_API_URL: vnc-monitoring.company.com | HTTP status code: 404
```

HTTP Status Codes:

- **200**: Success - API is working properly
- **401**: Unauthorized - Check credentials
- **404**: Not Found - Check URL/hostname
- **500**: Server Error - vNetC system issue
- **000**: Connection failed - Network/DNS issue

OPTION 8: SET DEBUG LOGGING

```
8. Set logging level to DEBUG
```

Purpose: Enables detailed logging for troubleshooting

Effect: Changes `LOG_LEVEL=DEBUG` in monitoring.ini

When to use: - Investigating specific issues - Before reproducing problems for support - Detailed system analysis

> ⚠ **Warning**
>
> Debug logging generates large log files and may impact performance

---

OPTION 9: POPULATE KNOWLEDGE GRAPH

```
9. Perform a ETL call to populate the knowledge graph
```

Purpose: Manually triggers data extraction, transformation, and loading

What it does: Executes the knowledge graph population process

When to use:

- After configuration changes
- To refresh Monitoring data
- Testing data pipeline functionality

Expected Behavior: Process runs and populates the knowledge graph with current data

---

Copyright BE Networks 2026

OPTION 10: SET ERROR LOGGING

```
10. Set logging level to ERROR
```

Purpose: Reduces logging to only error messages

Effect: Changes `LOG_LEVEL=ERROR` in monitoring.ini

When to use:

- After debugging is complete
- To reduce log file sizes
- For normal production operation

OPTION 11: CREATE DEBUG SNAPSHOT

```
11. Create Monitoring debug snapshot for Verity Support
```

Purpose: Generates comprehensive diagnostic package for technical support

Prerequisites:

- Must attempt steps 1-9 first
- Specify which step failed

What it includes:

- Complete system status check
- All container logs (latest 10,000 lines)
- Error logs for each container
- System resource information
- Docker container status and stats
- Configuration settings
- Time synchronization status
- API connectivity results

Output Files: - `monitoring_debug_package_YYYYMMDD_HHMMSS.tar` - Complete debug package

Archive Contents:

```
monitoring_service_status_results_YYYYMMDD_HHMMSS.txt
<container_name>_latest.log (for each container)
<container_name>_errors.log (for each container)
system_info.txt
```

When to use: When you need to provide comprehensive system information to Verity Support

---

OPTION 12: RESTORE DOCKER VOLUME

```
12. Restore Docker Volume
```

Purpose: Provides restoration of Docker volume.

Choose your docker volume from the options provided.

Options:

12.1. **Prometheus**: Restore Prometheus Volume from backup

Purpose: Restores the Prometheus Docker volume from a backup.

Copyright BE Networks 2026

What it does:

- Prompts the user to select one of the existing Prometheus backups
- Performs a Docker Volume backup based on the selected backup

When to use:

- After data corruption in Prometheus
- To revert Prometheus data to a previous state
- For disaster recovery

12.2. Quit

OPTION 13: REINSTALL OF LATEST MONITORING APPLICATION PACKAGE

```
13. Reinstall of Latest Monitoring Application Package
```

What it does: Reinstalls the latest Monitoring Application Package

OPTION 14: CREATE A PROMETHEUS BACKUP

```
14. Create a Prometheus Backup
```

What it does: Creates a Prometheus backup that can be imported via Option 12: Restore Docker Volume.

OPTION 15: EXIT

```
15. Exit
```

Exists the menu.

DAILY HEALTH CHECK

1. Run Option 1 (Check services)
2. Run Option 2 (Check time sync)
3. Run Option 7 (Check API)

TROUBLESHOOTING WORKFLOW

1. Set debug logging (Option 8)
2. Check all services (Option 1)
3. Verify connectivity (Options 6, 7)
4. Perform ETL test (Option 9)
5. Create debug snapshot (Option 11)
6. Reset logging level (Option 10)

PRE-SUPPORT CONTACT

1. Complete steps 1-9 in order
2. Note which step fails
3. Create debug snapshot (Option 11)
4. Provide the generated tar file to support

Expected Output Examples

HEALTHY SYSTEM OUTPUT

Copyright BE Networks 2026

```
Select:
    1. Check if the Monitoring services are running
    ...
Enter your choice: 1

ETL services are running.
Grafana services are running.
Prometheus services are running.
Promtail services are running.
Telegraf services are running.
Neo4j-graph services are running.
Loki services are running.
Cadvisor services are running.
Alertmanager services are running.
Node-exporter services are running.
AI-service services are running.
Neo4j-vector services are running.

---------------------------------------------
Total Services Tested: 12
Currently Running: 12
Not Running: 0
---------------------------------------------

Press Enter to continue...
```

PROBLEM SYSTEM OUTPUT

```
ETL services are running.
Grafana services is not running.
Prometheus services are running.
...
---------------------------------------------
Total Services Tested: 12
Currently Running: 11
Not Running: 1
---------------------------------------------
run the command 'sudo docker compose -f /be_monitoring/docker-compose.yml up -d' and re-run this script
```

Troubleshooting Common Issues

NO DOCKER CONTAINERS RUNNING

Symptoms: All services show "not running"

Solution:

```
sudo docker compose -f /be_monitoring/docker-compose.yml up -d
```

API CALLS FAILING (HTTP 401)

Symptoms: Option 7 shows HTTP status code 401

Cause: Invalid API credentials

Solution: Check `/be_monitoring/.env` file for correct credentials

API CALLS FAILING (HTTP 000)

Symptoms: Option 7 shows HTTP status code 000

Cause: Network connectivity issue

Solution: 1. Check Option 3 for correct FQDN 2. Verify network connectivity: `ping <fqdn>` 3. Check DNS resolution: `nslookup <fqdn>`

TIME SYNCHRONIZATION ISSUES

Symptoms: Option 2 shows `System clock synchronized: no`

Solution:

```
sudo systemctl restart systemd-timesyncd
sudo timedatectl set-ntp true
```

DEBUG SNAPSHOT CREATION FAILS

Symptoms: Option 11 fails to create archive

Copyright BE Networks 2026

Cause: Insufficient disk space or permissions

Solution:

```
# Check disk space
df -h
# Ensure proper permissions
sudo chown -R $USER:$USER ./
```

Log File Locations

SCRIPT OUTPUTS

- Debug snapshots: `./monitoring_debug_package_YYYYMMDD_HHMMSS.tar`
- Container logs: `./monitoring_logs_snapshot/`

SYSTEM LOGS

- Docker containers: `sudo docker logs <container_name>`
- System logs: `/var/log/syslog`
- Monitoring logs: `/be_monitoring/logs/` (if exists)

Best Practices

REGULAR MAINTENANCE

1. Run daily health checks (Options 1, 2, 7).
2. Monitor log levels (keep at ERROR for production).
3. Perform periodic ETL operations (Option 9).

BEFORE SUPPORT CONTACT

1. Always run the complete diagnostic sequence (1-9).
2. Create debug snapshot (Option 11).
3. Note exact error messages and timing.
4. Include system information and recent changes.

SAFETY GUIDELINES

1. Test changes in non-production first.
2. Always create debug snapshots before major changes.
3. Keep debug logging temporary (return to ERROR level).

Recovery Procedures

COMPLETE SYSTEM RECOVERY

If multiple services are down:

```
# Stop all containers
sudo docker compose -f /be_monitoring/docker-compose.yml down

# Start all containers
sudo docker compose -f /be_monitoring/docker-compose.yml up -d

# Wait 2 minutes, then check status
sleep 120
# Run troubleshooting script Option 1
```


Copyright BE Networks 2026

DATA RECOVERY

If data appears corrupted:

1. Use **Option 12** to restore the Prometheus volume from backup.
2. Re-run setup script if necessary.

Support Information

WHEN TO CONTACT SUPPORT

- Multiple services consistently failing
- API connectivity issues persist after basic troubleshooting
- Data corruption or loss
- Before using Option 999

INFORMATION TO PROVIDE

- Debug snapshot tar file from Option 11
- Specific error messages
- Recent system changes
- Output from Options 1, 2, 3, 4, 5, 7

SELF-HELP RESOURCES

- Check Docker container logs: `sudo docker logs <container_name>`
- Verify system resources: `free -h` and `df -h`
- Check network connectivity to vNetC host
- Review monitoring.ini configuration file

This troubleshooting script is designed to be comprehensive and self-contained, providing both diagnostic capabilities and maintenance functions for Monitoring.

Copyright BE Networks 2026

4.7.8 Change Sets

A **Changeset** is a logical grouping of pending changes that are applied to the system as a single change. Making changes to provisioning objects within a changeset ensures that multiple changes are compatible and consistent with the desired network state.

Related Information

- Verity API

Copyright BE Networks 2026

Export Import

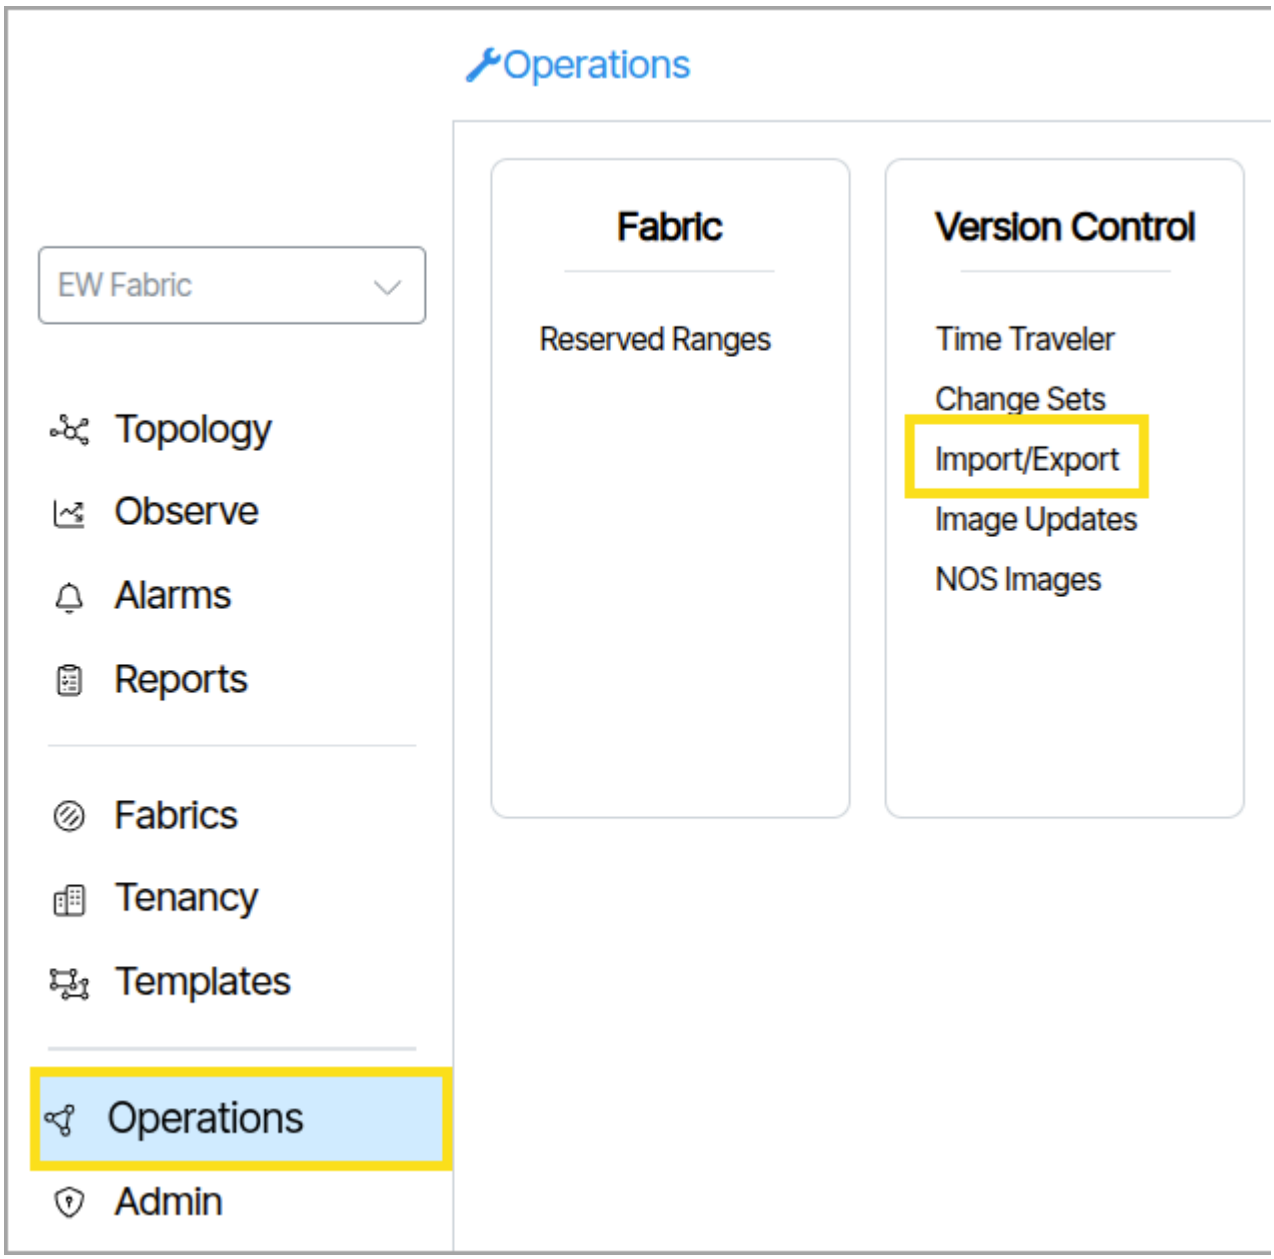
4.7.9 Import/Export

The **Import/Export** feature allows users to import or restore all provisioning settings back into the system, export or save all system provisioning to a text file, and compare provisioning by viewing differences between the current system provisioning and a text file before importing.

When to use CSV Import

CSV file imports can be used to create new systems from scratch. However, in most cases, it is easier to use the UI to create objects with fewer than 10 instances. The typical scenario where CSV imports are most efficient is when creating or modifying switch port provisioning or switch port labeling, especially when there are more than 64 ports per switch. For switch provisioning, it is recommended to first create examples of single port and LAG provisioning in the UI before exporting the files for updates.

To access the Import/Export feature go to **Operations** and select the text titled **Import/Export**.



Copyright BE Networks 2026



CSV Import Export Overview

Verity supports multiple methods for editing data objects. Examples include editing directly on the network map and within provisioning object blocks. Verity also offers bulk edit modes—both on the map using controls within each switch and by modifying provisioning data through reports. However, in cases such as creating new networks or implementing large-scale changes to an existing infrastructure, using the CSV import/export functions can provide a more efficient approach to bulk editing. CSV files can be used to create new systems from scratch or to manage updates by exporting the data, editing it offline, and re-importing it.

Verity models provisioning as a collection of managed objects. This collection, combined with live network topology discovery, is used by the system to generate the provisioning for the managed physical devices.

Files for the managed objects include:

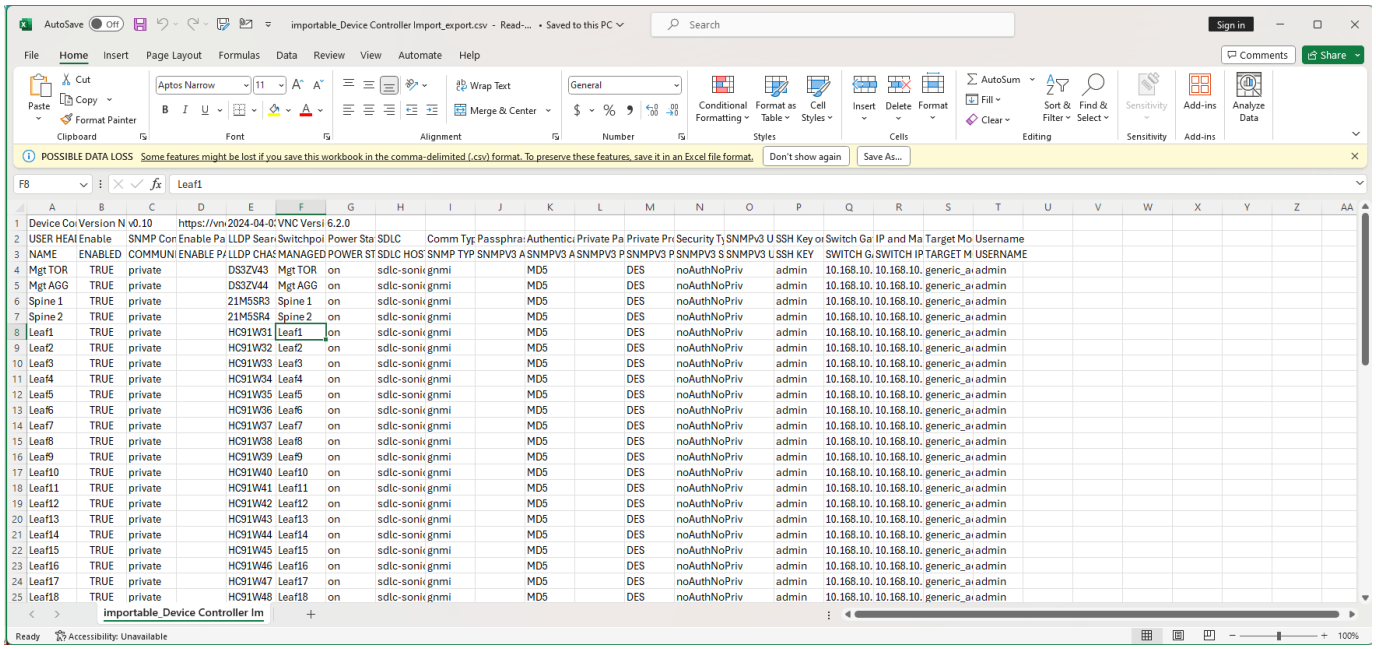
1. Tenants
2. Services
3. Ethernet Port Profiles
4. Ethernet Port Settings
5. Endpoint Bundles
6. Switchpoints (Endpoints)
7. Device Controllers
8. Gateways
9. Gateway Profiles
10. LAGs
11. Route Maps
12. IP Prefix Lists
13. Route Map Clauses

The files are database- and system-independent, making them easy to share between systems. Additionally, the import process is designed to support importing in any order by creating temporary references until the related managed objects are imported from the appropriate files. Each row in the file represents an instance of an object, and the columns represent the various parameters associated with that object, as shown below.

Switch provisioning can be found in the “Endpoint Bundles” CSV file, while switch labels and port enable settings are in the “Switchpoint” CSV file.

Example CSV files can be provided by customer service upon request, or you can create your own by exporting the files.

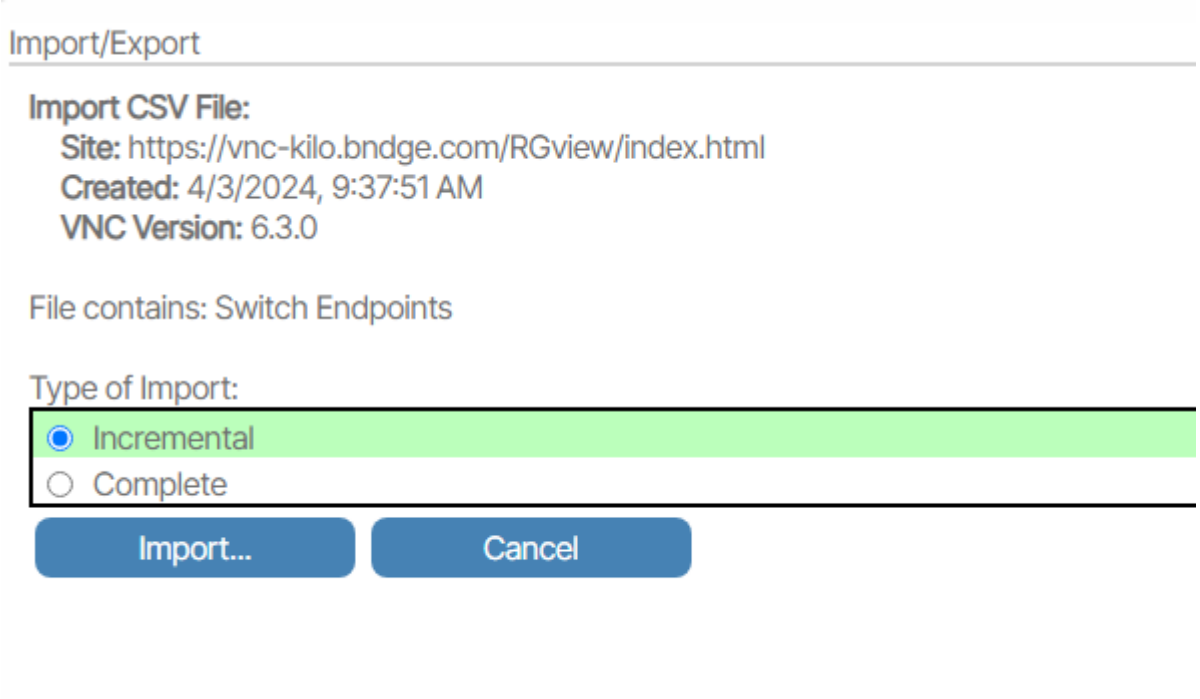
Copyright BE Networks 2026



Complete vs. Incremental

There are two types of CSV imports: Incremental and Complete. The Incremental import allows users to add only new items to the system, without affecting existing ones. In contrast, the Complete import assumes the file contains the full list of objects of a given type and deletes any existing objects of that type that are not included in the file. Incremental updates may include existing objects; if an object and all its parameters match an existing entry, it is ignored.

Using Incremental import is more common and generally safer, as it does not delete any objects from the system.



During the import process, the system performs basic validation of the content and then presents the user with a summary of the expected updates. These changes can be reviewed in detail before they are committed to the system.

Since it is generally good practice to save your current system state before performing any bulk updates, it is recommended to use the system's Time Traveler feature to create a complete provisioning snapshot beforehand.

Copyright BE Networks 2026

Compare

Compare is used to evaluate the differences between the current provisioning and a file prepared for import. To perform a comparison, click the **Compare** button. This feature displays the differences between the current provisioned configuration and the configuration defined in the imported file.

Importing the file after comparison will overwrite all existing provisioning objects.

Clear

Clear is used to remove devices and/or provisioning objects of various types from the system. When selected, all objects specified under the What to clear menu will be deleted. 

Parameter Names and Field Values

In some cases, the parameters and values used in the Verity UI *may not match* the column headers or values in the CSV file entries. Therefore, it is considered best practice to create one or two objects of the types you plan to import, then export the CSV files for those objects to use as a starting point for updates. The files can be renamed to any name, as long as they retain the CSV extension. Additionally, CSV file formats are specific to each Verity software release and should not be used directly between releases without first analyzing the format differences across versions.

Copyright BE Networks 2026

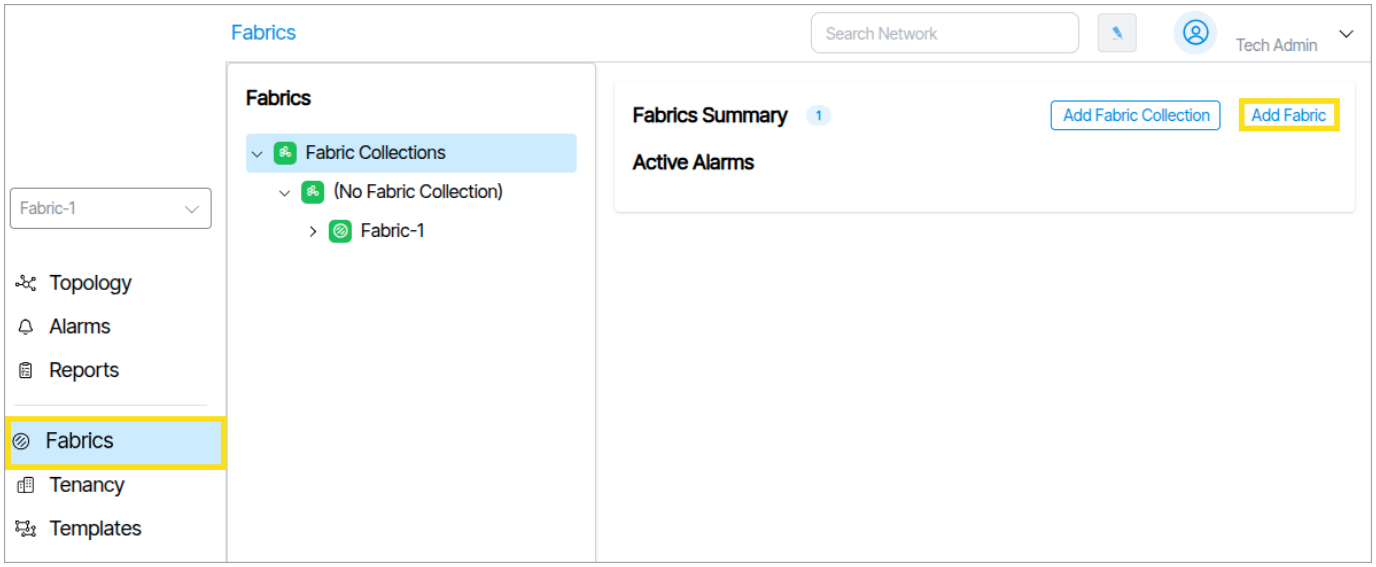
4.7.10 Spectrum X Fabric Device Import

Create the Fabric before importing device information. If you have already created the Fabric, jump to the section below titled **Import Device Information**.

Creating a Fabric

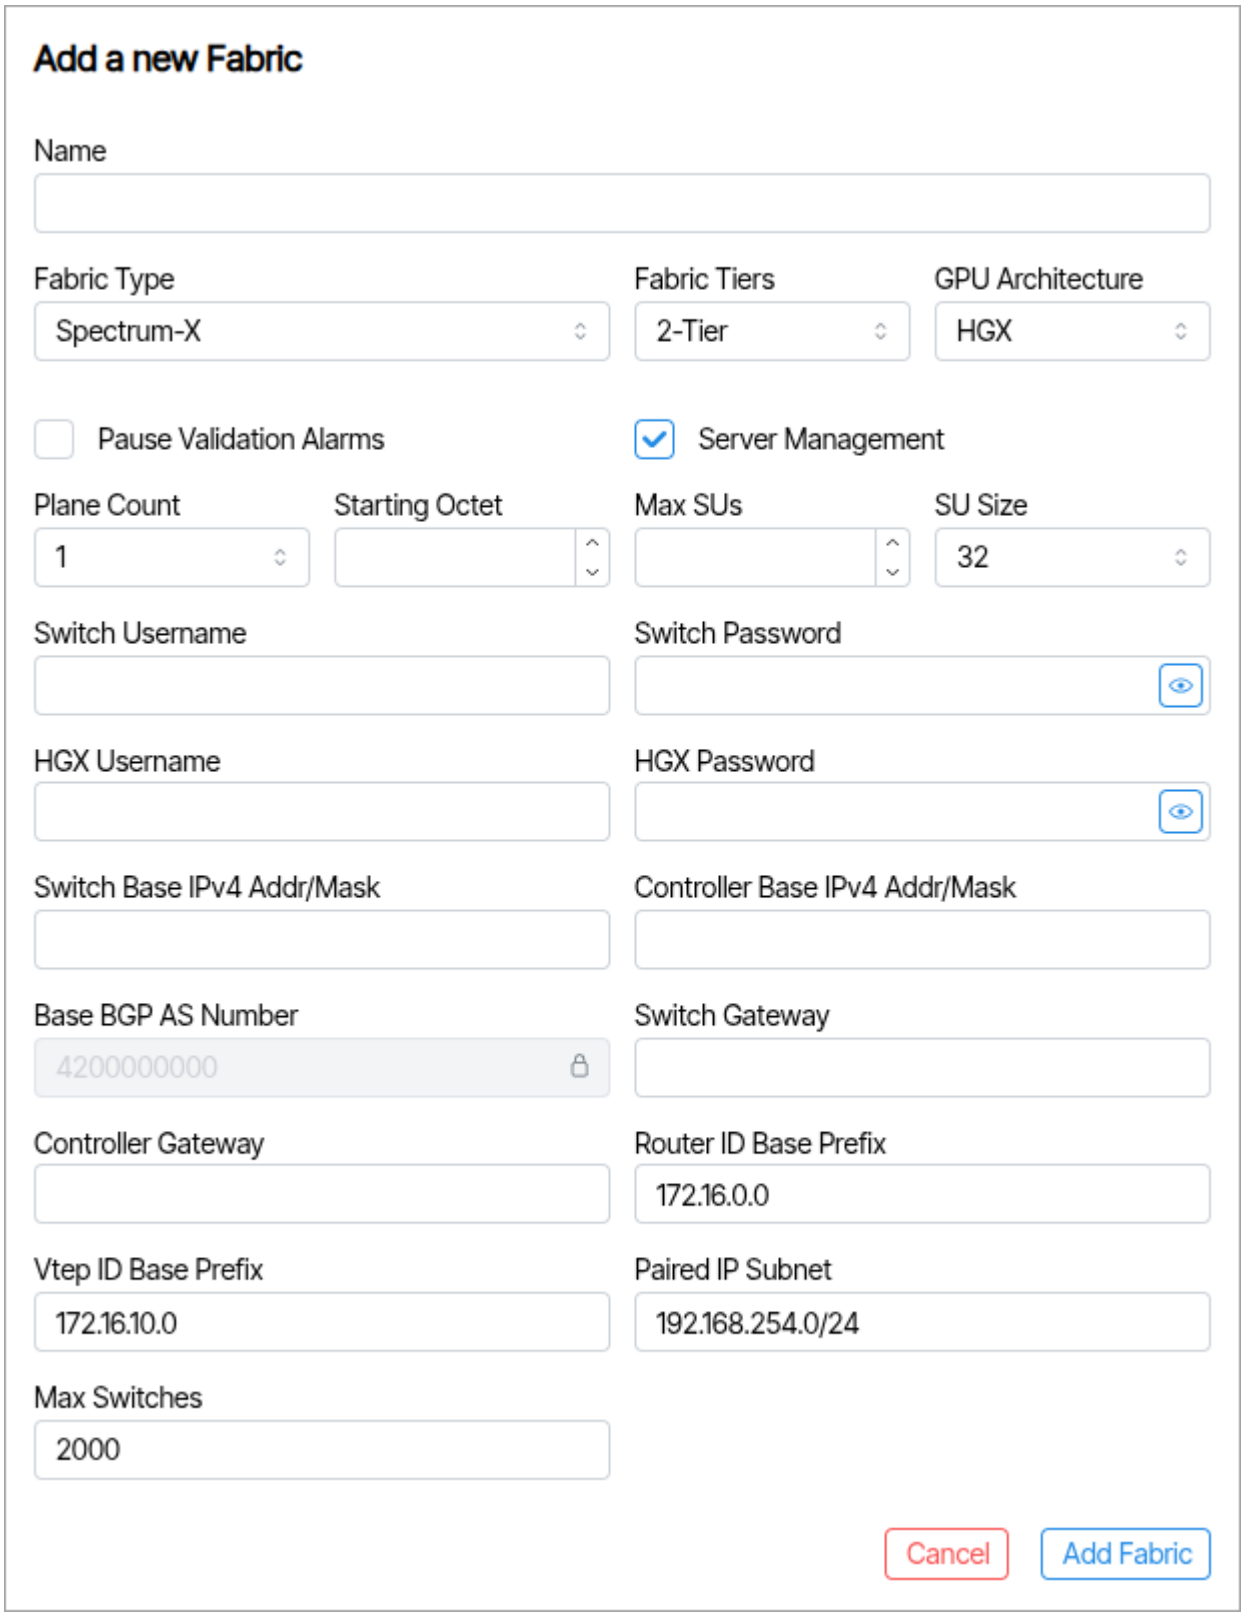
To create a new fabric:

1. Navigate to **Fabrics → Fabric Collection**.
2. Click the **Add Fabric** button. This will launch a form to select the Fabric type and enter details.



The **Fabric Type** needs to be set to a Spectrum X Tier Fabric. All other fields are required. When completed click **Add Fabric**.

Copyright BE Networks 2026



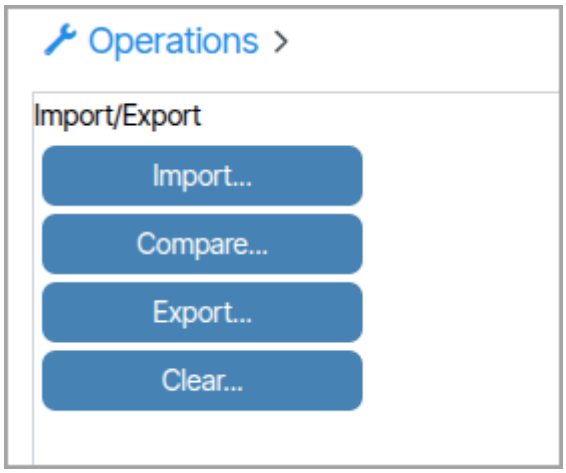
Import Device Information

After creating the Fabric, use the import file to bring in device information. The image below shows a sample device import file. The site name is PGA-Complex — in a real-world system, the fabric importing this file would share the same name.

A example of this file titled **FCA Verity Import Example File HGX** is available in Downloads

Copyright BE Networks 2026

To import the file go to **Operations → Import/Export** (). Click **Import** and select the file.



Copyright BE Networks 2026

4.7.11 Image Updates

Successful device software upgrades typically necessitate a device reboot. However, forcing a full network-wide restart for each device update is impractical and unacceptable. Image Update Sets solve this problem.

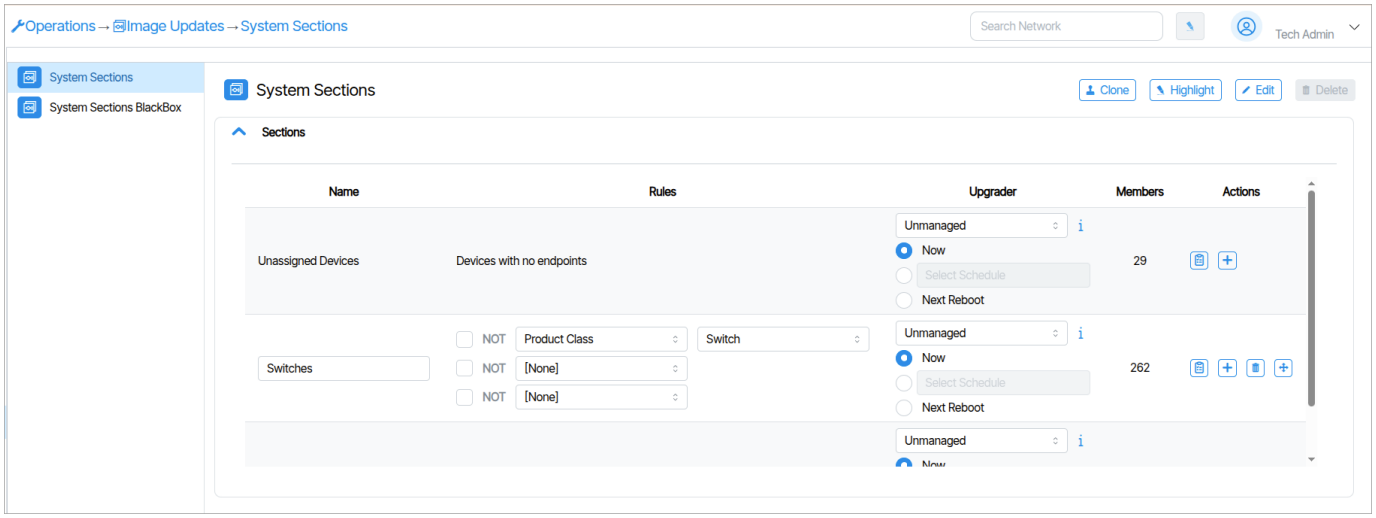
Image Update Sets allow you to organize devices into user-defined groups and target these distinct groups for software updates. This approach enables the system to leverage the redundancy of active switches to maintain network availability during the update process. This approach trades some capacity during the update, but it allows you to safely upgrade sections of your network without shutting down the entire system.

NOS Upgrades

UPGRADE BY IMAGE UPDATE SETS

> **Notice**
>
> This section describes the mechanics of using Image Update Sets; it does not cover how to apply them to any specific configuration or deployment scenario.

To access Image Update Sets, go to **Operations/Image Updates**.



The presented window is composed of rows and columns, with each row representing an Image Update Set and each column representing information specific to that Image Update Set. The **Members** field displays how many devices are in the Image Update Set. The columns titled **Name**, **Rules** and **Upgrader** are user defined values. The **Actions** column contains controls to create/delete new Image Update Sets as well as a button to open a Report listing devices that match the Image Update Set rules.

Name

Name is where you name your Image Update Set.

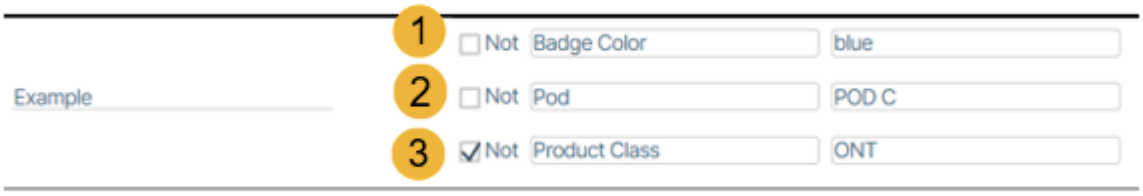
Rules

Devices are grouped into Image Update Sets by using conditional logic in the **Rules** column. A Rule contains a **Not** clause, a condition (named **Rule Type**), and a field containing values contingent on the condition (named **Rule Value**). The Rule value options change depending on the selected Rule type. All three fields are shown in the following image.

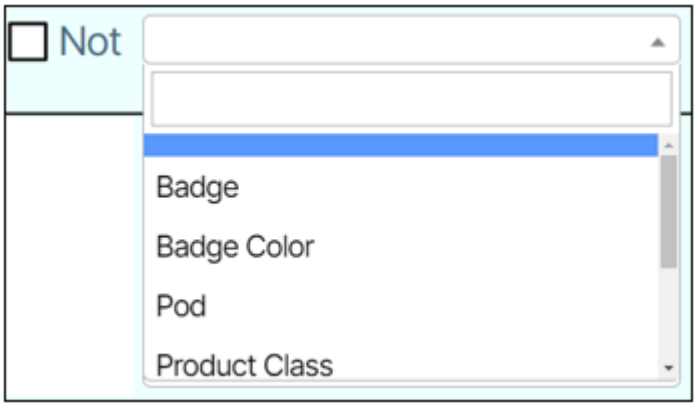


Each Image Update Set can have up to three Rules.

Copyright BE Networks 2026



There are many conditions to choose from. Clicking the conditional field provides a scrollable list.

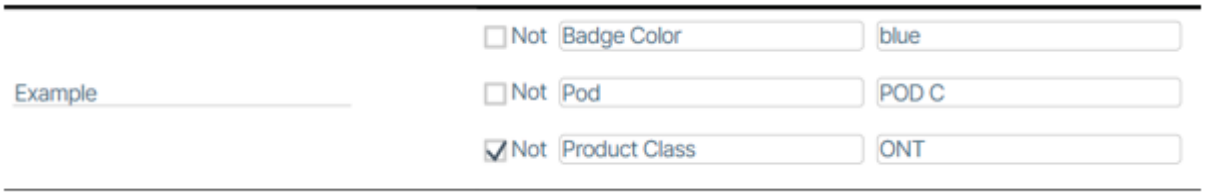


Rule Example

The example in the following image selects all devices assigned a **Badge Color** of **blue**.



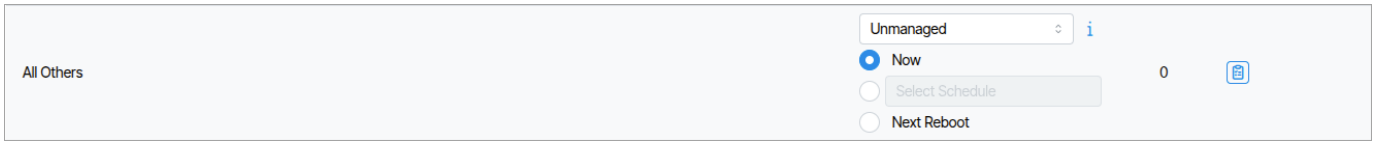
This more complex example selects all devices in the network with a blue **Badge Color** that are in **POD C** and are **Not** of the product class ONT.



Order Matters

When writing Image Update Sets, it is important to understand that order matters. If a device is included in an Image Update Set, it cannot be included in subsequent Image Update Sets further down the page.

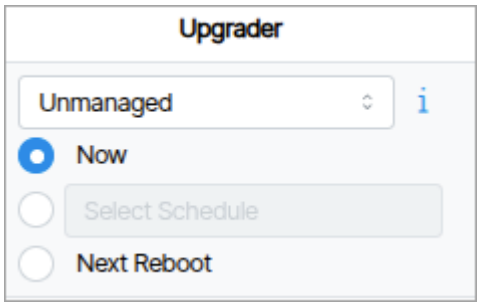
Any devices not explicitly assigned to an Image Update Set are assigned to **All Others** on the bottom row.



Copyright BE Networks 2026

UPGRADER

The Upgrader column is where you select the software version package and schedule the reboot time. This action is applied to all devices that fall within the Image Update Set Rules.



The first field specifies the software version, followed by three options for selecting the device's restart time.

Now

This performs the image transfer and reboots immediately.

Month/Day/Year, Hour: Minute: Second AM/PM

When you select this option, the image file is transferred and written to the standby partition of the device. The device then reboots at the scheduled time.

> **Notice**
>
> Switch Devices have active and standby partitions. This allows you to load device image software during normal operation and reboot at a scheduled time.

Next Reboot

When you select this option, the image file is transferred and written to the standby partition of the device and updates on the next *manual* reboot.

MEMBERS

Members shows the device count for the Image Update Set. Double-clicking area opens a Report listing all devices in the set.

How to Upgrade by Device

Each device can be upgraded individually independent of Image Update Set. To do so, zoom into the device **SW Version** tile. 

In the window that appears you have the choice to choose **Target Package** by **Set** and **By Device**. Select **By Device**. The drop-down menu lists the available upgrade packages. Choose the one you desire. Save the choice by clicking the save button.

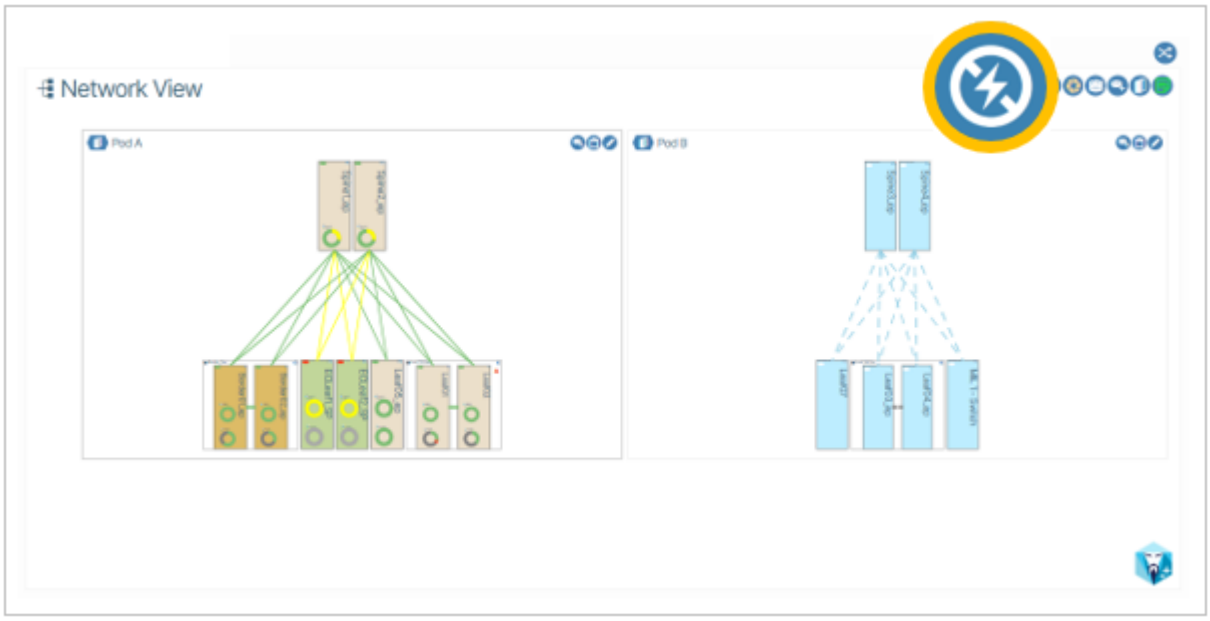
> **Warning**
>
> The moment you save, the device will reboot!

Copyright BE Networks 2026



How to Enable/Disable Devices from Receiving Software Updates.

In Network View, the **Image Upgrades of all Devices Enable/Disable** button determines if a system will accept software updates.

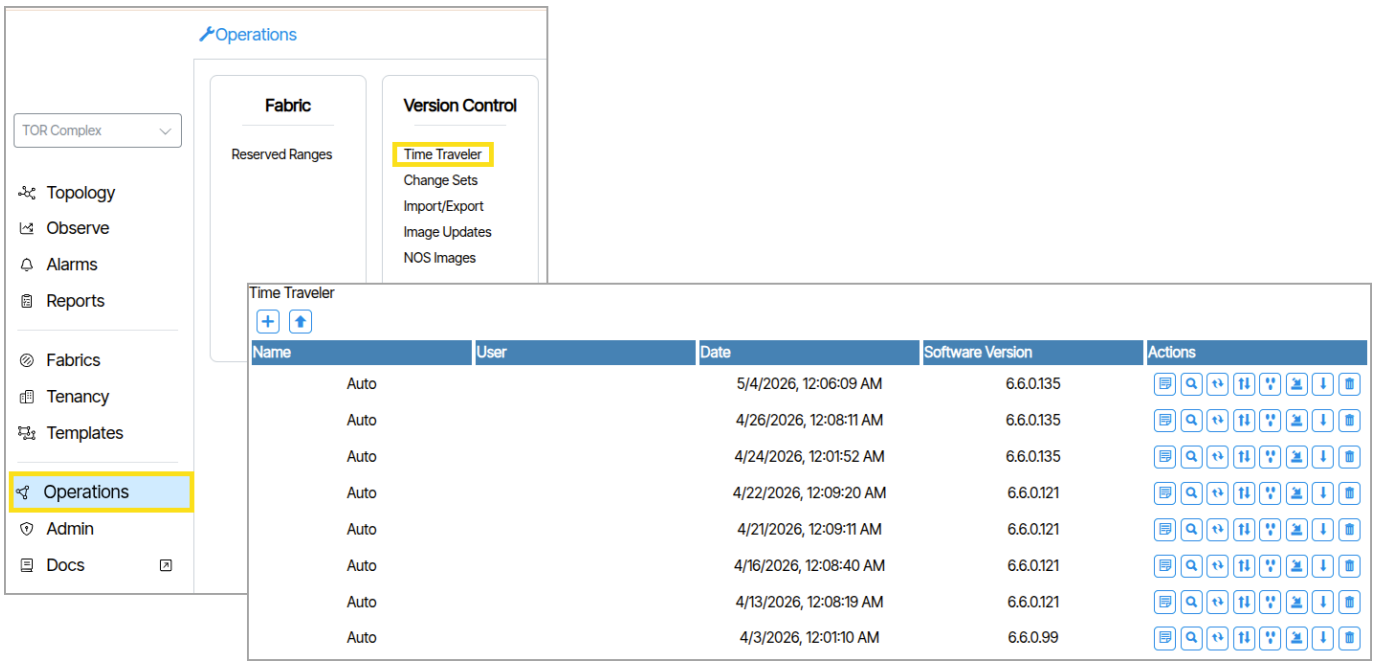


Copyright BE Networks 2026

Backup

4.7.12 Network Time Traveler

Network Time Traveler is a tool to create and restore the intended configured state of the network. Time Traveler exercises version control on the configured objects in Verity. Verity derives device configurations from these objects, therefore the system can recalculate the entire network configuration when administrators use the Time Traveler function. Verity does not backup the actual device configurations.



Manually Create Backup

1. Navigate to **Operations/Time Traveler**.
2. Click the **Create a New Network Time Traveler Backup** button to create a new backup. 

Revert State Using Time Traveler

1. Navigate to **Operations/Time Traveler**.
2. Determine the the backup to revert to (Select by date).
3. Click the corresponding **Restore Network Time Traveler Backup** icon. ()

Copyright BE Networks 2026

TIME TRAVELER ACTIONS



1. Write and/or View Note of backup
2. Compare to Current
3. Compare to Previous Backup
4. View Changed Provisioning Objects Since Backup
5. View affect Ports Since Backup
6. Restore Network Time Traveler Backup
7. Download Network Time Traveler Backup
8. Delete Network Time Traveler Backup

AUTOMATIC TIME TRAVELER BACKUPS

Time Traveler backups are automatically created once a day if there is a change in the network detected. A maximum of 10 is retained by default and the oldest one is removed if a new one is made.

Copyright BE Networks 2026

4.7.13 Reserved Ranges

> **Warning**
>
> If your system's network identifiers require custom ranges, it is strongly recommended to configure these settings *immediately after* installing the Verity software and *before* beginning the preprovisioning process. The preprovisioning process involves adding devices using the system initialization file and/or manually adding devices. Changing the network identifier ranges later may require reprovisioning the entire network.

The **Reserved Ranges** are a set of editable value ranges for networking identifiers used in the provisioning process. On a new system, it populates with working default values, so the user doesn't need to perform any custom configuration to have a working system.

To tailor the network setup to specific requirements, you may need to customize the allocated ranges of network identifiers. Customizing these settings allows for more precise control over network configurations and ensures compatibility with your network environment.

System Reserved Ranges

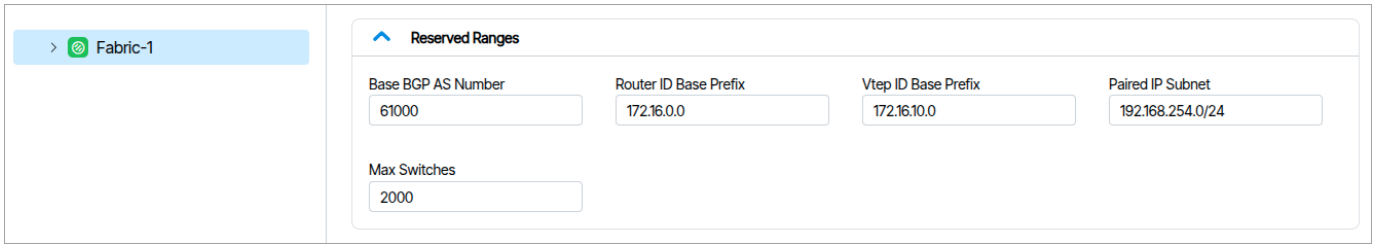
System Reserved Ranges apply to the entire system and are configured at **Admin → Reserved Ranges**.

Fabric Reserved Ranges

Fabrics → *Your Chosen Fabric* → Reserved Ranges

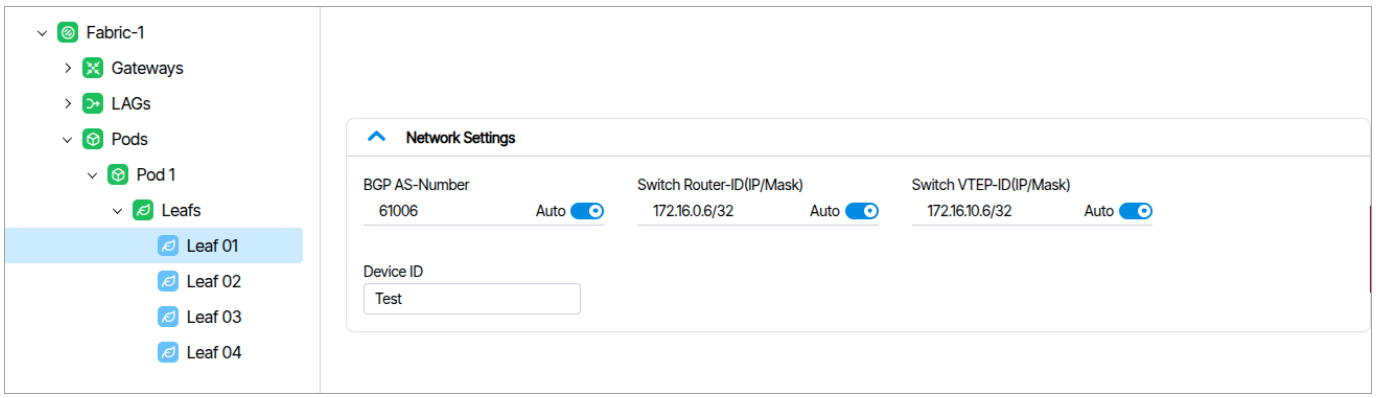
These Reserved Ranges are **Fabric** *specific* and may include the following:

- Base BGP AS Number
- Router ID Base Prefix
- VTEP ID Base Prefix
- Paired IP Subnet
- Max Switches



Device Values

Individual device values marked with the **Auto** checkbox inherit their value from the Fabric's reserve range settings.



Copyright BE Networks 2026

4.7.14 Badges

Badges group devices visually. Administrators assign each badge a name, a color, and a number.



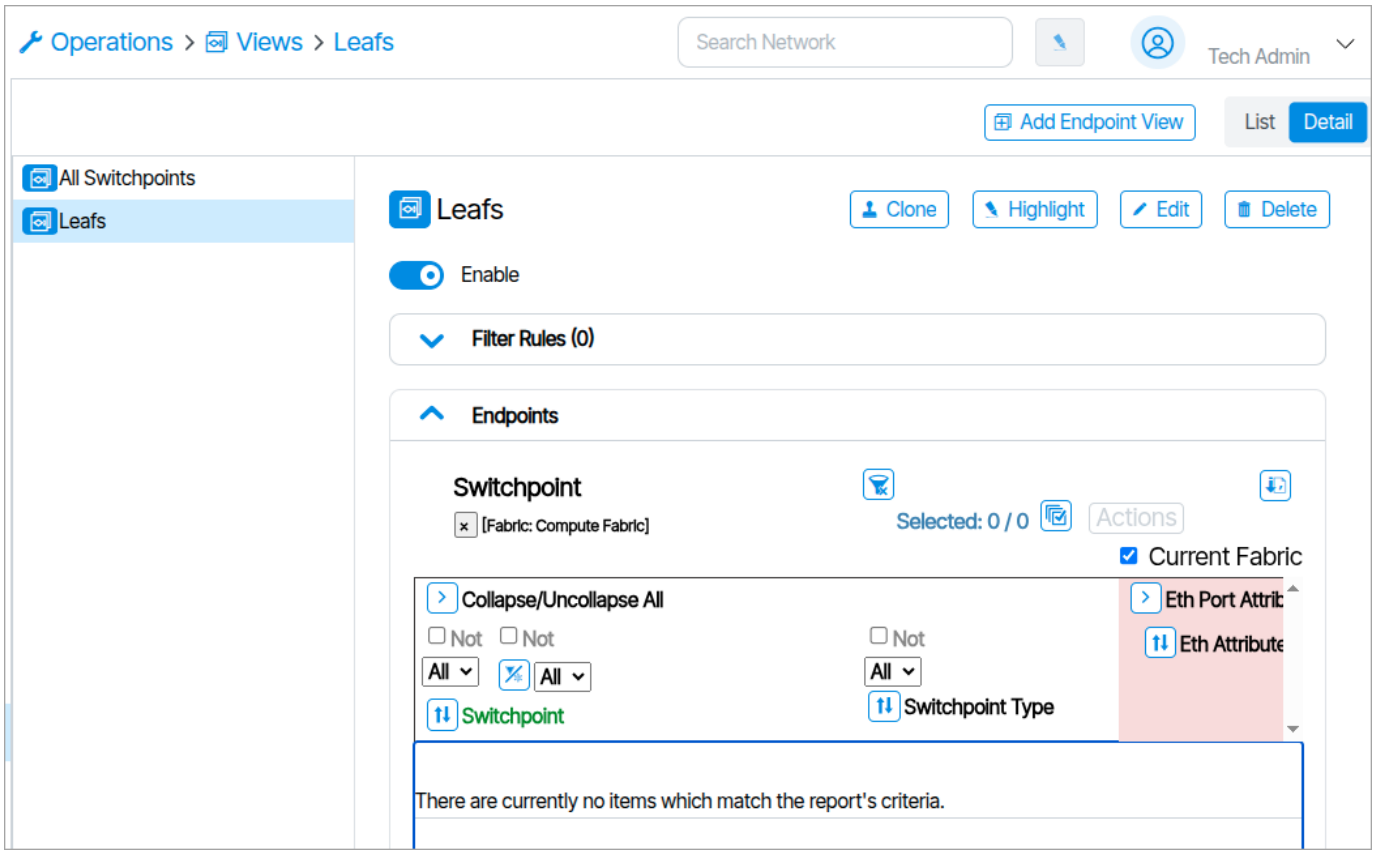
Creating and Assigning Badges

1. Double-click **Topology** and click **Network View**. Enable focus of a selected device.
2. Click the **Open for Edit** button (or double-click the Badges field). 
3. Click set badges icon. 
4. Enable the check-box next to the phrase Create New Badge.
5. Set a Name, Color and Num for the badge.
6. Click Yes to create the badge.

Copyright BE Networks 2026

4.7.15 Views

Views provide graph and table presentations of all Switchpoint data. The information includes comm and provisioning status. All Switchpoints appear by default. Administrators can also create custom Views that group together selected Switchpoints.



Copyright BE Networks 2026

4.8 Administration

Licensing RBAC System

4.8.1 Admin

The **Admin** section provides high-level system configuration, user management, security, operational management tools, and settings.

VNFs

The **VNFs** section manages virtual machine installations of the vNetC and SDLC, including the ACS, Backend Service(s), and AI Service.

VNETC COMMANDER

This section is composed of important processes working within the vNetC virtual machine.

System Applications

System Applications are:

- ACS
- DHCP Server
- Backend Service(s)
- AI Service

Installing System Applications

1. Go to **Admin**.
2. Select **VNFs**.
3. Click **Add a System Application** button ()
4. In the window that appears, click the desired System Application ()

The example below demonstrates how to install Backend Service(s); the steps are applicable to the other System Applications.

How to Install or Update Backend Service(s)

Backend Service(s) is a System Application that increases UI performance. To add **Backend Service(s)** to your application do the following:

1. Go to **Admin/VNFs** and navigate to the section of the window titled **System Applications** .
2. Click the create button ()  .
3. Choose **Add a Backend Service(s)** from the prompt ().
4. In the window that appears, enable the application by clicking the enable button ()and fill out the fields with the relevant IP information.
5. Save your work by clicking the checkbox icon().
6. The yellow **Status** box text will read **Awaiting Status**. When it changes to **Connected**, the process is complete. ()

LICENSING

Licensing displays the license expiration date, support contact information, and reports on license and physical port usage.

Copyright BE Networks 2026

RESERVED RANGES

The **Reserved Ranges** are a set of editable value ranges for networking identifiers used in the provisioning process. On a new system, it populates with system-wide working default values, so the user doesn't need to perform any custom configuration to have a working system.

Read about Reserved Ranges Per Fabric here..

AI PROVIDER CONFIG

These setting allow you to choose your AI provider.

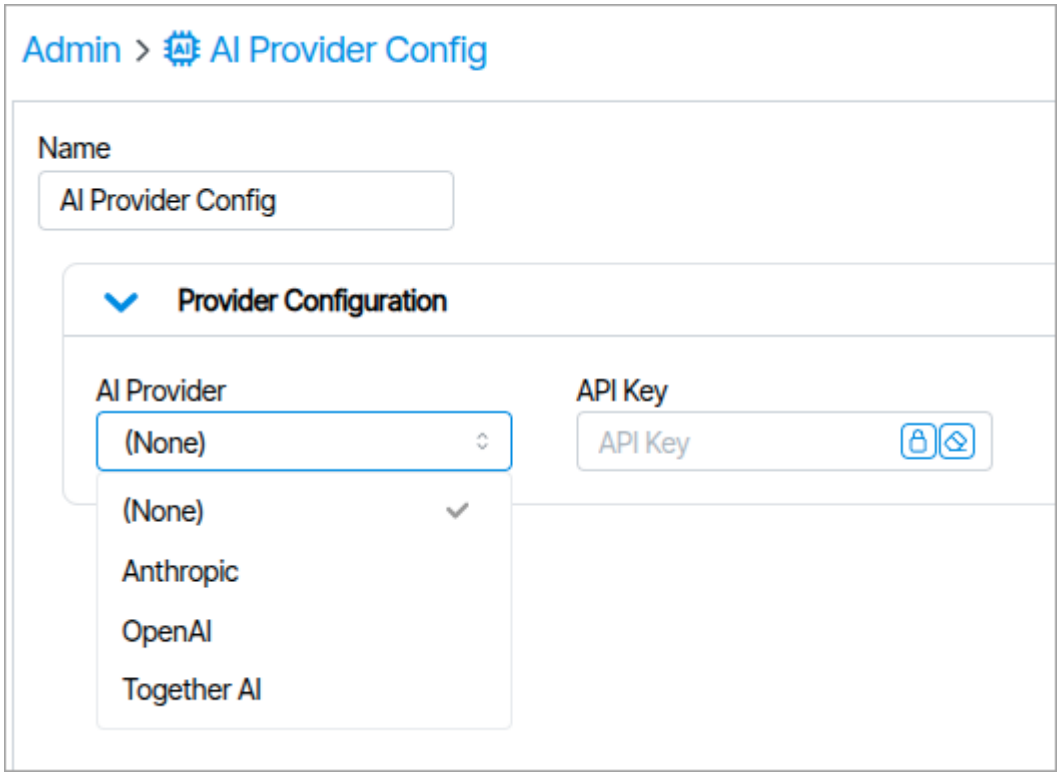
How to Configure AI Ops

AI Ops is Verity's AI-powered messaging assistant. To add **AI Ops** to your application, set the **AI Provider Config** values by doing the following:

1. Go to **Admin → AI Provider Config**.
2. In the **Provider Configuration** form, set the **AI Provider**.
3. In the **API Key** field, click the **Edit Encrypted Field** icon and, in the form that appears, paste your API key.
4. Click **Save**.

How to Disable AI Ops

1. Go to **Admin → AI Provider Config**.
2. In the **Provider Configuration** form, set the **AI Provider** to **None**.



Users

Users provides tools for authentication, user settings, and role assignment.

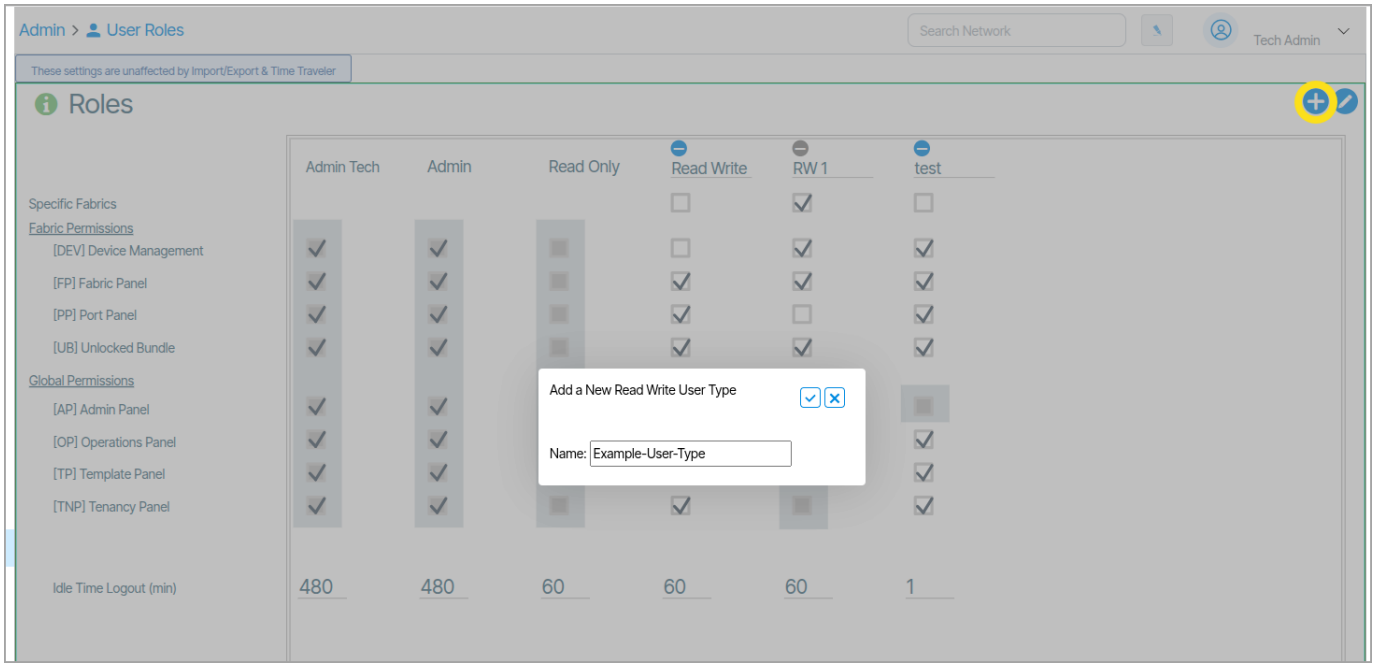
ROLES

Verity supports role based access (RBAC) permissions scheme to partition the various workflows to operational personnel.

Copyright BE Networks 2026

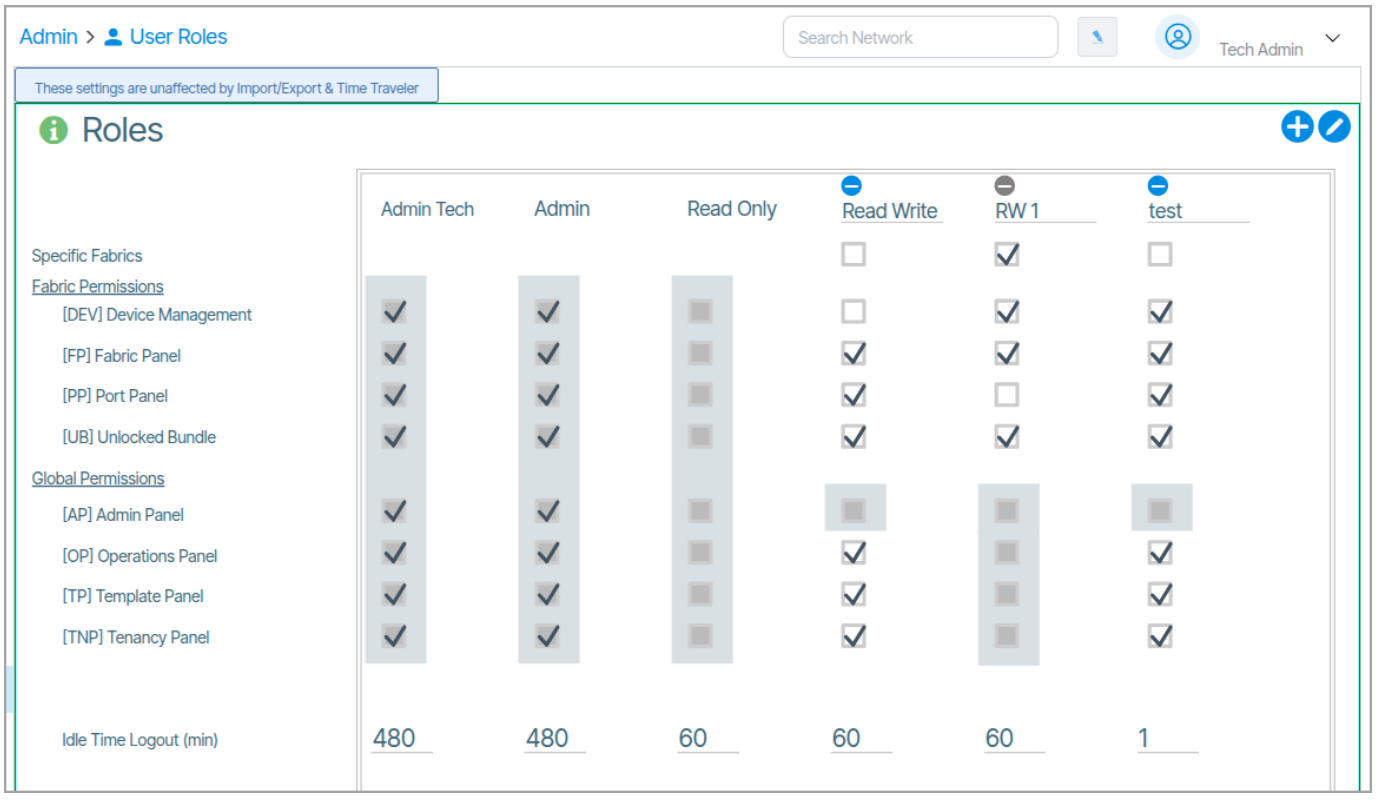
The **Roles** settings are found under **Admin → User Roles**. Use them to grant or remove editing access to specific features for each role (also called a **User Type**).

User Types are created by clicking the **Create a New User Type** button, creating a name for the User Type and submitting the form.



The **Roles** window contains a collection of checkboxes with each **User Type** (role) listed as a column item and each Permission is listed as a row item. You enable and disable checkboxes to determine what Permissions are accessible to each Role.

- **User Types **(roles) have titles like Admin, Admin Tech and Read Only.
- **Permissions** have titles like Fabric Panel, Admin Panel and Tenancy Panel.



Copyright BE Networks 2026

4.8.2 User Permission Roles

Permission	Parameters	Department / Team
Specific Fabrics	Limits the Fabric Permissions below to the fabrics assigned in **Admin > User Fabric Permissions**.	*Scope setting (not a role)*
Fabric Permissions		
[DEV] Device Management	Capability to manage Devices.	**Device Operations**
[FP] Fabric Panel	Capability to edit Fabrics, PODS, Planes, fabric-level objects and both types of Bundles (locked and unlocked)	**Fabric Architecture**
[UB] Unlock Bundle	Edit unlocked Bundles	**Service Provisioning**
Global Permissions		
[AP] Admin	Grants access to edit all features in the GUI	**Super Administrator**
[OP] Operations Panel	Capability to Import/Export. Manage Badges, Views and Time Traveler	**Change & Release Management**
[TP] Templates Panel	Edit Tenants, Bundles and Provisioning objects	**Network Engineering**
[TNP] Tenancy Panel	Edit Tenants, Services and Tenant Gateways	**Tenant & Routing Services**
Idle Time Logout (min)	Number of minutes a user can remain inactive before they are logged out	**Any**

Specific Fabrics

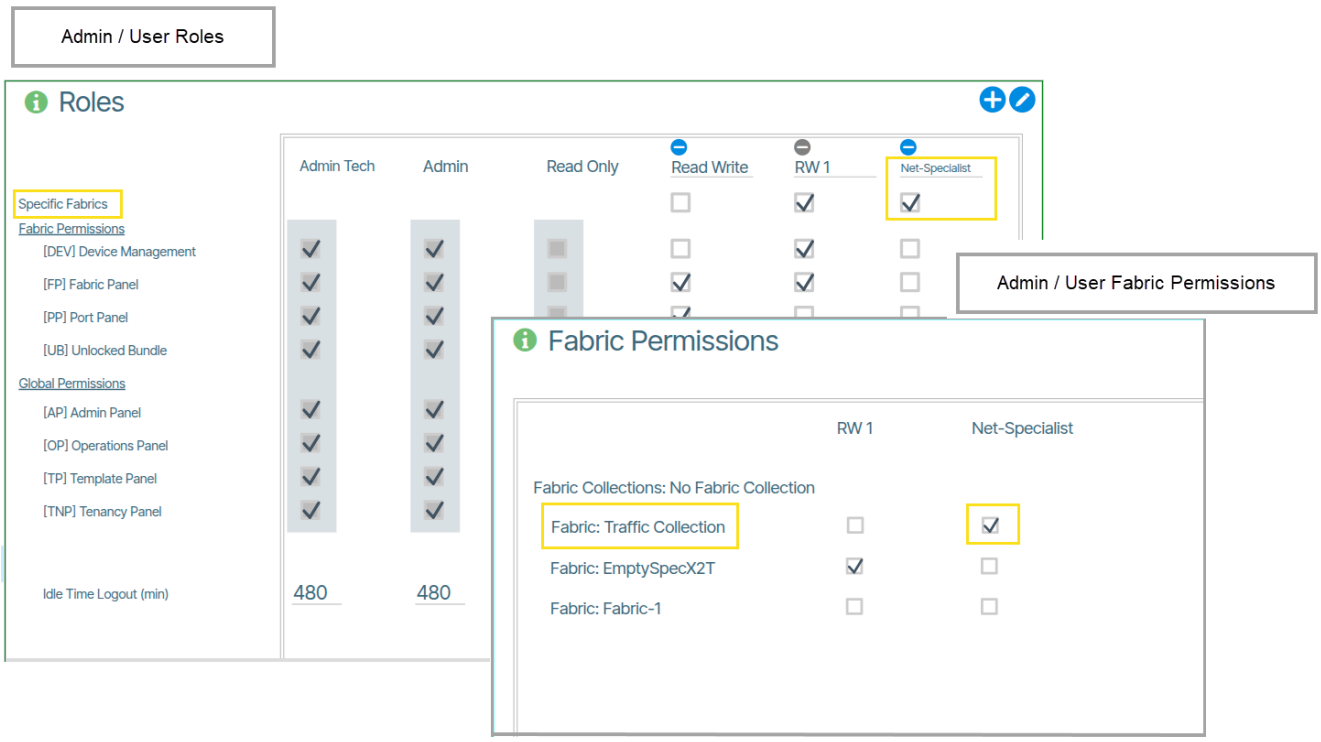
When **Specific Fabrics** is enabled on a User Type, users with that type can access only the Fabrics an administrator assigns to them in **Admin → User Fabric Permissions**.

USER FABRIC PERMISSIONS

Admin → User Fabric Permissions is where administrators assign individual **Fabrics** to **User Types**.

In the image below, a **User Type** titled **Net Specialist** is granted editing privileges for a Fabric named **Traffic Collection**.

Copyright BE Networks 2026



HOW FABRIC PERMISSIONS WORK

In a User Type, **Fabric Permissions** and **Specific Fabrics** work together:

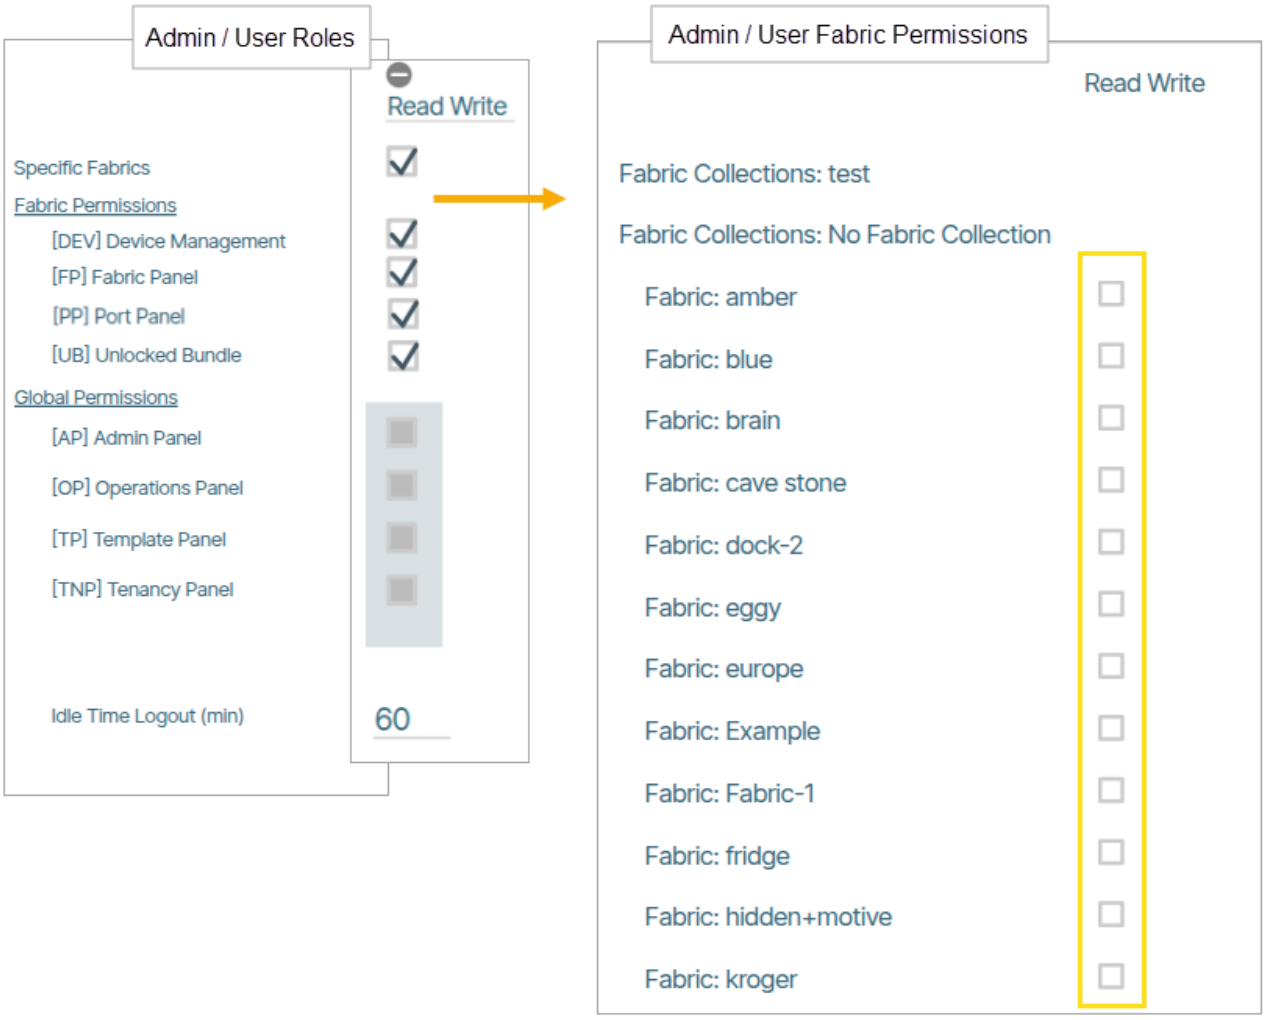
- **Fabric Permissions**: *what* the user can edit (the features you check).
- **Specific Fabrics**: *where* they can edit (every fabric, or only the fabrics you pick).

Specific Fabrics	Result
Unchecked	The user can edit the selected features on **all fabrics**.
Checked	The user can edit the selected features **only on the fabrics assigned to them** in **Admin > User Fabric Permissions**.

Note: Selecting **Specific Fabrics** alone does not grant any access. You must also select at least one feature under **Fabric Permissions**; otherwise, the user cannot edit anything on their assigned fabrics.

The Read/Write example below meets this requirement. However, if no fabrics are selected in **Admin → User Fabric Permissions**, the Read/Write role behaves as Read Only.

Copyright BE Networks 2026



CONSIDERATIONS

- **Specific Fabrics needs Fabric Permissions.** Checking Specific Fabrics grants nothing by itself, the User Role must select at least one Fabric Permissions option.
- **Operations is nearly as powerful as Admin.** Time Traveler and Import/Export can roll back or replace the configuration for the whole network. NOS image updates can reboot many switches at once. Give this panel only to trusted staff.
- **Authentication settings are in two panels.** RADIUS is in Operations. TACACS, LDAP and Device AAA Profiles are in Templates. Anyone with either panel can change how users log in to the switches.
- **Templates and Tenancy depend on each other.** Route maps are created in Templates and assigned in Tenancy. A Tenancy-only user can attach route maps but can't create or edit them.

Copyright BE Networks 2026

USER ACCOUNTS

Assigning Users to User Types (roles)

In **Admin → User Accounts**, each user is assigned a User Type. The User Type is under the heading titled **Access Level**.

1. From the **Admin** window select the **Users Accounts** item box. 
2. In the window that appears, click the Access Level menu item twice for the chosen user.  You can also select the edit icon () in the upper right corner of the **User Accounts** window to edit the page.
3. Choose the Username you want to assign and set the **Access Level** to the desired User Type (role). 
4. To complete the process, click the checkmark () in the upper right corner of the **User Accounts** window to save your work.

USER RADIUS LOGINS

Remote Authentication Dial-In User Service (**RADIUS**) is a networking protocol that provides centralized Authentication, Authorization, and Accounting (AAA or Triple A) management for users who connect and use a network service. RADIUS is a client/server protocol that runs in the application layer and can use either TCP or UDP as transport. The devices associated with Endpoints control access to a network as they contain a RADIUS client component that communicates with the RADIUS server. RADIUS is often the back end of choice for 802.1X authentication as well.

To access the RADIUS Logins settings:

1. Navigate to **Admin**.
2. In the Users column select **User RADIUS Logins**. 

RADIUS LOGIN SECURE DEPLOYMENTS

The Verity platform provides the option for RADIUS integration iVN-Authorization/Server IP/Port/Shared secret are required inputs. User permission level will be passed down according to the following levels:

Vendor IDs are assigned by IANA

| **VENDOR** | **iPhotonix** | **49683** |

RADIUS dictionary for Verity remote authentication provides the initial authorization level. Once authenticated, accounts are automatically created and updated with the iVN-Auth level. The "superuser" user does not actually exist within Verity, it is mapped to (same as) the "admin" user. The "authenticated" value indicates that the user is authenticated but that the authorization is based on internal records as managed by the Verity Admin page. The initial level is readonly. An authentication returned with no iVN-Auth entry will be ignored. Only the first iVN-Auth will be honored.

ATTRIBUTE	iVN-Auth	String
VALUE	iVN-Auth	superuser
VALUE	iVN-Auth	admin
VALUE	iVN-Auth	readwrite
VALUE	iVN-Auth	readonly
VALUE	iVN-Auth	authenticated

Optional authorization modifier This instructs the Verity Web UI to modify the operations and results available to the user. For example, the value "tech" modifies the web behavior for "admin" and "superusers" to include more technically-oriented features.

ATTRIBUTE	iVN-Auth-Modifier	string
VALUE	iVN-Auth-Modifier	tech

Copyright BE Networks 2026

USER LIMITS

User Limits sets the number of simultaneous users that can access the website simultaneously

USER SSH KEYS

Here you can view, upload and delete user SSH keys.

Certificates

Certificates manages the import and revocation of application certificates. This is where you work with SSL certificates, certificate chains and certificate revocations.

IMPORTING CERTIFICATES

To import a Certificate:

1. Navigate to the **Admin** section and select **Certificates** ().
2. Double-click the tile that describes the certificate you want to upload.
3. Either drag the certificate to the drag and drop section or click the browse button and select your certificate file to upload it. 

VNETC SERVER CERTIFICATE

For the **vNetC Server Certificate** panel, certificate files must be in PEM format. The **vNetC Server Certificate** should include

- Private Key File
- Certificate File
- Optional CA Intermediate Certificate

These should all be in PEM format, concatenated in that order. The CA Root Certificate should not be included.

Once uploaded, the **vNetC Server Certificate** is automatically added to the backend web path. Although this does reconfigure the backend process, this is done so that new connections will be handled using the new certificate but old connections will continue the operate as before, so there is no outage of the web service during a certificate update.

> **Information**
>
> Verity can be configured to support client certificate authentication. This operates in a "lockdown" mode that meets US DoD requirements. It does impose strict access requirements on all transports in the systems and demands client certificates for all users, devices, and external application access.

Software

The **Software** section controls software lifecycle management including vNetC platform packages, hardware-specific firmware images, bundled image packages, configuration file management, and application package deployment.

Branding

Branding customizes the user interface with white-label options, including the login banner, top navigation logo, browser favicon, splash screen, lower-right copyright text, the product documentation hyperlink destination and general application settings.

Network

Network configures system-wide network behavior, management connectivity, WAN routing, and optional feature toggles.

Copyright BE Networks 2026

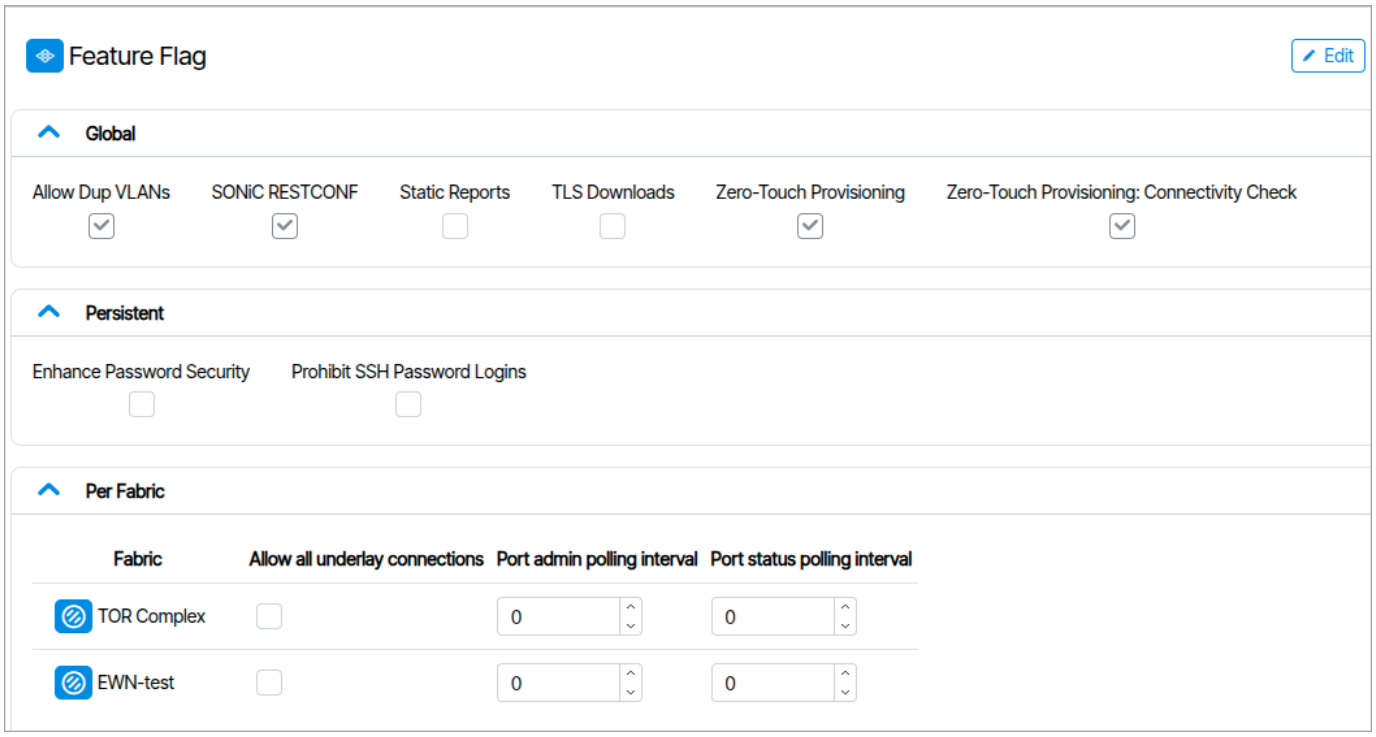
ADMIN SETTINGS

Admin Settings is accessed from Admin/Admin Settings and lets you configure the following:

- vNetc Address on Management VLAN
- Permissible IP Address Ranges on Managed Devices
- Customized Download Address/FQDN

FEATURE FLAG

Feature flags are used to enable or disable visibility of options in the UI and other system wide behaviors. Please consult with customer service regarding the usage of these settings beyond the defaults that are set upon installation.



Global

Allow Dup VLANs

Services to allow duplicate VLNs

SONiC RESTCONF

Enable or disable SONiC RESTCONF support

Static Reports

Used for extremely large fabrics.Stops reports from dynamically updating unless requesting.

TLS Downloads external Use HTTPS instead of HTTP for firmware downloads.This causes extra vNetC load.

Zero-Touch Provisioning

Generates ZTP files.

Zero-Touch Provisioning: Connectivity Check

Copyright BE Networks 2026

When unchecked, skips the post-provisioning connectivity check for in-band managed switches, where the management path may not be immediately reachable after ZTP completes.

Persistent

Enhanced Password Security Enable or disable password constraints and expiration. Existing users are not affected until a password change takes place.

Prohibit SSH Password Logins

- Disables password-based logins to the vNetC, Monitoring, and controller command-line prompts.
- SSH access to the vNetC remains enabled, but only via keys registered under **Admin → User SSH Keys**.
- SSH access to Monitoring and the controllers is permitted from the vNetC using `ns_cli`.
- Restrictions may take up to an hour to apply to the Monitoring and controller targets.
- If the change is required immediately, restart the target.

Per Fabric

This section contains three fields per Fabric.

- **Allow all underlay connections**

Allows underlay connection between PODS

- **Port admin polling interval**

Polling intervals values in seconds, set if aggressive reporting is not enabled.

- **Port status polling interval**

Polling intervals values in seconds, set if aggressive reporting is not enabled.

ROUTING

Routing includes the following WAN IP information.

- Hostname (FQDN)
- RAT Port Range
- WAN IP Source (DHCP or Static)
- WAN Default Route
- WAN IP Address
- DNS Server

External

External integrates with external monitoring and management systems and automates database backup configurations.

SNMP

Configure SNMP trap forwarding to external monitoring systems.

SYSLOG

Configure Syslog forwarding to external logging servers.

DB BACKUPS

How to Automate DB Backups

1. Select **Admin/DB Backup** from the Navigation Bar. 

Copyright BE Networks 2026

2. Enter the required information:
   - **BWLimit(kbps)**: Bandwidth Limit
   - **User**: Username on the remove server
   - **Server**: Remote server name (can be IP address or hostname)
   - **Port**: IP port of the remote server
   - **Filepath**: Path on the remote server
3. Enable the **Enabled** and **Host Subdir** boxes.
4. Once Enabled, click the **Auth Key** button to copy the authorization key. 
5. In the **backup server**, add the authorization key to the system's authorized keys store, eg: **.ssh/authorized_keys**.
6. Save.

Plugins

DELL OME INTEGRATION

System connects to Dell OME servers to collect hardware statistics from Dell servers and display them in dashboards.

Copyright BE Networks 2026

4.8.3 NTP Administration

Introduction

NTP (Network Time Protocol) keeps system clocks synchronized within a few milliseconds of Coordinated Universal Time (UTC). This ensures certificate expiration times are handled accurately as well as viewing unified system logs.

NTP can operate in one of three modes:

1. Unauthenticated
2. Authenticated via a shared key
3. Authenticated via PKI (using a variant of NTP called NTS)

Out of the box, the vNetC uses unauthenticated NTP with the "Frankenbox" host and a public pool of NTP servers. If no "Frankenbox" host is present, that server entry never synchronizes, and the public servers maintain the time instead.

Configure a Local Unauthenticated NTP Source Clock

Use this when the vNetC has no internet access and your NTP server does not require authentication.

1. Copy `/usr/ns/etc/ntp-lead-servers.conf` to `/var/ns/etc/ntp-local.conf`
2. Edit `/var/ns/etc/ntp-local.conf` to add or remove NTP servers
3. Edit `/etc/ntp.conf` to uncomment the line that references `ntp-local.conf` and comment out the `ntp-lead-servers.conf` line
4. Run `service ntpd onerestart` (an error message saying `WARNING: failed to start ntpd` can be ignored)
5. Verify with `ntpq -p` that your server appears in the output

Configure Authenticated NTP (Shared Key)

Use this when your environment requires that the vNetC authenticate its time sources — for example, in JITC-compliant or air-gapped deployments. Your NTP server administrator will need to provide the shared keys.

STEP 1: CREATE THE KEY FILE

With shared-key authentication, the servers and the clients must all have a key in common, though it is not necessary for all servers and all clients to have the same keys.

One easy way to handle this is to get a copy of the server key file, or you can construct the file by hand. The file should be placed at `/var/ns/etc/ntp.keys` and should look like this:

```
# ntpkey_MD5key_ivn-ext2.dev.ipho.lan.3898159044

# Wed Jul 12 08:57:24 2023

# Sample keys only!
# Your keys will be provided by your NTP server administrator
#
1 MD5 L";Nw<*.I<{4U0)247"i  # MD5 key
2 MD5 &>10%XXK90'51VwV<xq~  # MD5 key
3 MD5 1b4zLW~d^!K:]RsD'qb6  # MD5 key
4 SHA1 a27872d3030a902568446c751b4551a7629af65c  # SHA1 key
5 SHA1 21bc3b4865dbb9e920902abdccb3e04ff97a5e74  # SHA1 key
6 SHA1 267736fe24fef5ba85ae11594132ab5d6f6daba9  # SHA1 key
```

The key index numbers (1 through 6 above) must match between the server and client configurations. Some organizations issue per-network keys to limit exposure — if that is the case, you will need to edit the file manually rather than copying it directly from the server.

STEP 2: CREATE THE LOCAL CONFIGURATION FILE

Create `/var/ns/etc/ntp-local.conf` with the following structure:

Copyright BE Networks 2026

```
server 192.168.0.41 iburst key 1
server 192.168.0.42 iburst key 2

keys /var/ns/etc/ntp.keys
trustedkey 1 2

# Optional: fall back to local hardware clock if all upstream servers are unreachable
server 127.127.1.1
fudge 127.127.1.1 stratum 10
```

Replace the IP addresses with your actual NTP server addresses. Using IP addresses rather than hostnames means time synchronization will still work if DNS is unavailable. The `key` value on each server line must match the index for that server in your key file. The `trustedkey` line must list every key index referenced in the server lines.

The `server 127.127.1.1` block is optional. It causes the vNetC to fall back to its local hardware clock if all upstream servers become unreachable, so that any downstream clients remain synchronized even if the time is no longer UTC-accurate.

STEP 3: APPLY THE CONFIGURATION

Edit `/etc/ntp.conf` to comment out the default servers line and uncomment the local config line:

```
includefile /usr/ns/etc/ntp-lead-head.conf
#includefile /usr/ns/etc/ntp-lead-servers.conf
includefile /var/ns/etc/ntp-local.conf
```

The NTP daemon restarts automatically when this file is saved.

STEP 4: VERIFY SYNCHRONIZATION

Run `ntpq -p` after a few seconds (allow up to a few minutes). A healthy output looks like this:

```
| remote | refid | st | t | when | poll | reach | delay | offset | jitter |
|---|---|---|---|---|---|---|---|---|---|
| \*192.168.0.41 | 166.129.200.1 | 4 | u | 324 | 1024 | 377 | 3.247 | -1.121 | 1.955 |
| +192.168.0.42 | 144.90.5.123 | 3 | u | 589 | 1024 | 377 | 2.503 | -0.820 | 1.581 |
| LOCAL(0) | .LOCL. | 10 | l | - | 64 | 0 | 0.000 | 0.000 | 0.000 |
```

A `*` in the first column means that server is actively selected for synchronization. A `+` means it is a candidate if the primary becomes unavailable. `LOCAL(0)` reflects the optional hardware clock fallback — because it is forced to stratum 10, it will only be used if all other servers become unreachable.

TROUBLESHOOTING

If `ntpq -p` shows `.INIT` in the `refid` column after several minutes, the most common causes are:

- The server hostname is not resolving — try using an IP address to rule out DNS issues
- The key on the client does not match the key configured on the server
- UDP port 123 is blocked between the vNetC and the NTP server

Copyright BE Networks 2026

4.8.4 High Availability (vNetC & SDLC)

Verity supports an automatic failover solution which provides support for high availability (HA) in the orchestration platform.

The solution requires two host systems. As these are operated in active/standby mode, each system must be sized to handle the full system load according to the resource calculator referenced in the installation instructions.

The feature is supported for the KVM hypervisor or VMware.

The optional Monitoring (Observability) VM is also involved in HA however switchovers do not occur based on failures of the Monitoring VM. Monitoring VM is installed in both server instances and are both active. Both instances receive statistics from the active vNETC. In the case where one Monitoring VM is offline for a long time (i.e. where significant statistics are lost) a mechanism is provided to copy the statistics database from backup or from the other Monitoring instance.

This section includes installation information as well as considerations for Verity upgrades while systems are in HA mode.

Installation Overview

Each system is installed per the standard instructions in the Quickstart section of this documentation. This section contains a configuration worksheet that should be filled out ahead of time with the overall network plan and address assignments in the environment containing the two host machines.

The process contains the following steps:

1. Review the contents of this complete section
2. Complete the Configuration Worksheet
3. Install each system per standard instructions. It is important that the steps of upgrading vNETC and SDLC to the latest release are done per instructions before proceeding the next step.
4. Once the full HA system is operational, and the system is planned to be using Monitoring, install Monitoring VM on side 1 first and complete the process, and then install Monitoring in side 2.

> **Important**
>
> The sections below titled "Creating the vNETCs" and "Creating the SDLCs" have specific information to guide the standard installation process.

> **Important**
>
> Since both SDLCs are configured with the same addresses they should not be allowed to connect to the same network until the HA process is completed on both systems. You can achieve this by always making sure one SDLC is shutdown or by having the second server completely isolated until HA configuration is completed.

> **Important**
>
> Both Monitoring instances must be installed pointing to the HA floating IP address only.

1. Configure vNetC HA mode

System Block Diagrams

Verity supports both single external (server) NIC and two external NIC configurations for installation. Please refer to those diagrams for the installation within the Quikstart section as your base starting point for HA.

Copyright BE Networks 2026

When using HA, the two installation options become either a two external NIC or three external NIC installations, respectively. The additional external NIC is used for the cross-connect communications between the two servers. The user may chose to use bond groups which increases the number of external NICs in their servers for any interfaces, but this is transparent to the Verity vNETC.

In normal operating mode, the active vNETC periodically transfers important state to the inactive vNETC in addition to continuously replicating POSTGRES database updates.

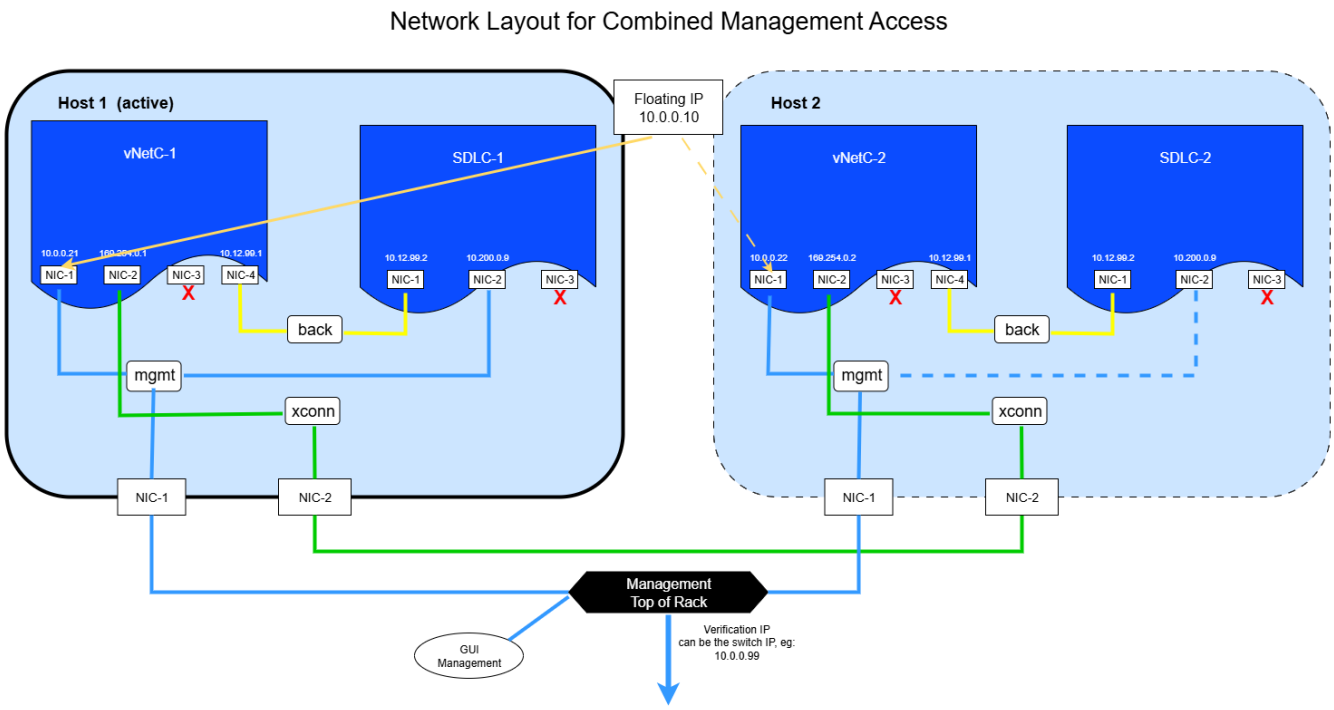
The additional state includes:

1. System configuration
2. Master key vault
3. Statistics data

The management switch architecture is up to the user (e.g. single switch, MLAG etc.) and is also transparent to the vNETCs.

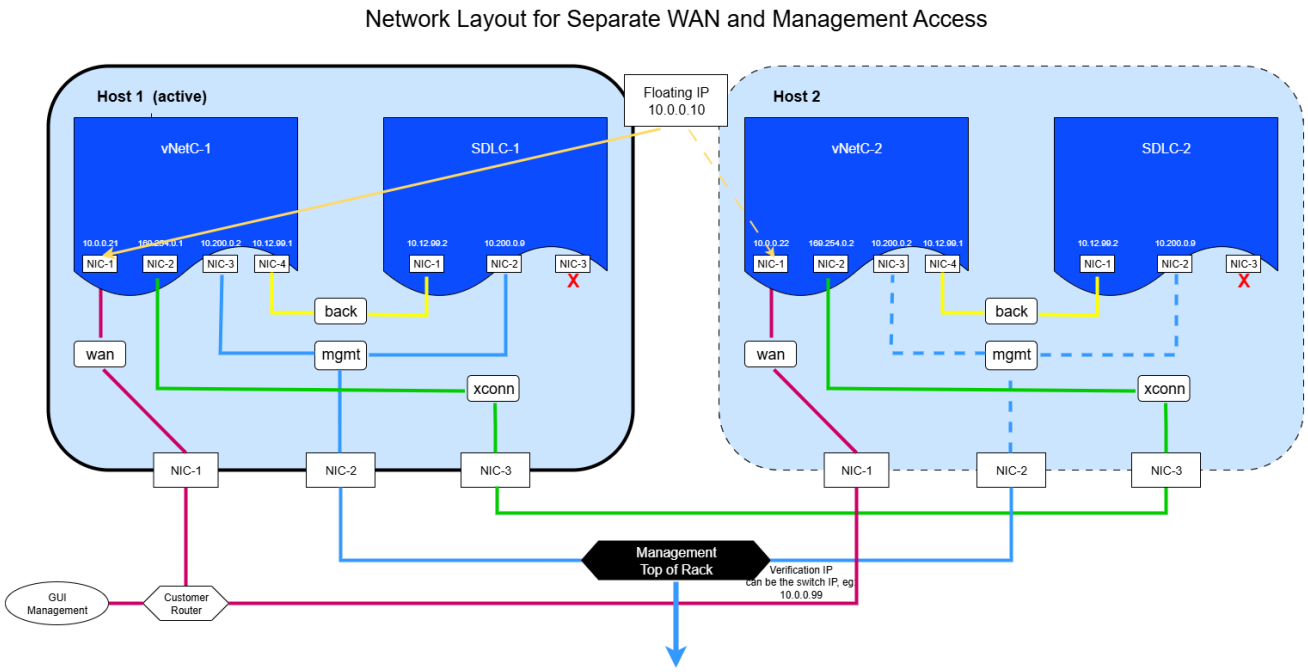
The HA system level diagrams for both installation options is shown below (the Monitoring VM is excluded from the diagram. It is simply connected to the mgmt interface in both use cases):

TWO EXTERNAL NIC BLOCK DIAGRAM



Copyright BE Networks 2026

THREE EXTERNAL NIC BLOCK DIAGRAM



Configuration Worksheet

Token	Description	Example	Your Value
MGMT_NET	The management network address	10.0.0.0	
WAN_NET	The WAN network address (can be the same as MGMT_NET)	10.0.0.0	
WAN_GW	The gateway address for the WAN network	10.0.0.1/24	
FLOATING_IP	The IP address that will access the Web GUI and float between redundant nodes	10.0.0.10	
FQDN	The FQDN for the floating IP (Needed if GUI will be accessed using a domain name)	vnetc.domain.com	
BACK_NET	The backdoor network used to manage the SDLC	10.12.99.0	
XCONN_NET	The cross-connect network address	169.254.0.0	
VERIFICATION_IP	The IP address of a switch management port that responds to ICMP echo-request	10.0.0.99/24	
DNS 1	The IP address of the primary DNS server	8.8.8.8	
DNS 2	The IP address of the secondary DNS server		

Copyright BE Networks 2026

COMMON SDLC ELEMENTS

Token	Description	Example	Your Value
SDLC_IP	SDLC IP address on management network (same for both SDLCs)	10.0.0.9	
SDLC_SN	SDLC serial number (same for both SDLCs)	SDLC-A	
SDLC_HOSTNAME	SDLC hostname (same for both SDLCs, can be the same as SDLC_SN)	SDLC-A	
ACS_IP	ACS IP address on management network (same for both SDLCs)	10.0.0.3	

KVM HOST ELEMENTS (NOT FOR VMWARE)

Token	Description	Example	Your Value
HOST1_IP	The IP address for accessing the first KVM server host on the management or WAN network	10.0.0.11/24	
HOST2_IP	The IP address for accessing the second KVM server host on the management or WAN network	10.0.0.12/24	

NODE 1 SPECIFIC ELEMENTS

Token	Description	Example	Your Value
NAME1	The hostname for first vNetC	vnetc1.domain.com	
VNETC1_IP	The IP address of the first vNetC in the WAN_NET or MGMT_NET network	10.0.0.21/24	
VNETC1_BACK_IP	The IP address of the first vNetC in BACK_NET	10.12.99.1/30	
SDLC1_BACK_IP	The IP address of the first SDLC in BACK_NET	10.12.99.2/30	
VNETC1_XCONN_IP	The IP address of the first vNetC in XCONN_NET	169.254.0.1/30	

NODE 2 SPECIFIC ELEMENTS

Token	Description	Example	Your Value
NAME2	The hostname for second vNetC	vnetc2.domain.com	
VNETC2_IP	The IP address of the second vNetC in the WAN_NET or MGMT_NET network	10.0.0.22/24	
VNETC2_BACK_IP	The IP address of the second vNetC in BACK_NET (can be the same as VNETC1_BACK_IP)	10.12.99.1/30	
SDLC2_BACK_IP	The IP address of the second SDLC in BACK_NET (can be the same as SDLC1_BACK_IP)	10.12.99.2/30	
VNETC2_XCONN_IP	The IP address of the second vNetC in XCONN_NET	169.254.0.2/30	

Time synchronization

The High Availability system requires that clocks on the two vNetCs are synchronized. This is set up automatically but will work only if the systems have NTP access to the internet. If not, NTP must be configured on the vNetCs to operate with an accessible NTP server on the local network.

Important

The following sections cover VMware and KVM specific information. Use the section appropriate for your installation.

Copyright BE Networks 2026

VMware Specific Configuration

VIRTUAL SWITCH AND PORT GROUP CONFIGURATION

1. Create a virtual switch called "xconn" with the physical port assigned that is used to connect the two servers with a direct connection.
2. Create a portgroup called "xconn" and select the "xconn" virtual switch. This is used for the heartbeat test between the servers.
3. Create a virtual switch called "back" with no external physical port.
4. Create a portgroup called "back" to use between the vNetC and SDLC.

CREATING THE VNETCS

Follow the Verity quickstart instructions to install the vNetC under VMware

Before powering up the vNetC, add a network adapter and connect it to the "back" port group.

Skip to the section following the KVM Specific Configuration section called "Configure vNetCs for HA mode"

KVM Specific Configuration

HOST OS INSTALLATION

The host should be installed with the Ubuntu-LTS server edition, with sshd enabled.

```
apt update; apt upgrade -y; apt autoremove -y
apt install -y net-tools iputils-ping rsync curl wget traceroute htop
apt install -y qemu-kvm libvirt-daemon-system libvirt-clients bridge-utils smartmontools
```

ADDRESSING AND BRIDGES

The host should have at least two physical NICs which we'll call "eno1" and "eno2".

"eno1" should be connected to the WAN network, and "eno2" should be connected to the "eno2" interface on the redundant pair.

If there is a management network separate from the WAN network, then it should be connected to a third physical NIC which we'll call "eno3".

The WAN network should be accessible to administrators who will use the vNetC GUI. This will have:

- two fixed addresses, one for each vNetC VM
- one floating address, used to access the active vNetC GUI

The management network should be configured with an IP space large enough for all managed devices to have an IPv4 address, plus:

- one fixed address shared by the two SDLC VMs
- one fixed address shared by the two ACS containers

The "eno2" interface is presented to the vNetC VM via a bridge "xconn". It should have a /30 interface (or larger). As this is a point-to-point connection, we typically use a link-local address like "169.254.0.1/30".

Each vNetC needs to be able to manage the SDLC running on the local host. This is done using a bridge called "back". This should be configured with an IP space large enough for a vNetC and each SDLC, so a /30 is the minimum size, but it makes sense to use a larger block to allow for future growth. In these examples we use "10.12.99.1/24".

In these examples, the netplan format is used, but it is also possible to use KVM network configurations. The components (vNetCs and SDLCs) are all configured with static addresses, so there is no need to provide DHCP support on any of the bridges.

NETPLAN CONFIGURATION EXAMPLES

Netplan example with management network and WAN combined. In this case, the third interface is not connected and not used but is needed to keep a consistent interface order.

```
network:
  ethernets:
    eno1:
      dhcp4: false
      dhcp6: false
    eno2:
```


Copyright BE Networks 2026

```
      dhcp4: false
      dhcp6: false
  bridges:
    wan:
      interfaces: [eno1]
      addresses:   [HOST1_IP]
      dhcp4: false
      dhcp6: false
      routes:
      - to: default
        via:   WAN_GW
    xconn:
      interfaces: [eno2]
      dhcp4: false
      dhcp6: false
    mgmt:
      interfaces: []
      dhcp4: false
      dhcp6: false
    back:
      interfaces: []
      dhcp4: false
      dhcp6: false
  version: 2
```

Netplan example with separate management network and WAN. In this case, the third interface is connected to the management port (eno3). Differences from above are highlighted in green.

```
network:
  ethernets:
    eno1:
      dhcp4: false
      dhcp6: false
    eno2:
      dhcp4: false
      dhcp6: false
    eno3:
      dhcp4: false
      dhcp6: false
  bridges:
    wan:
      interfaces: [eno1]
      addresses:  [HOST1_IP]
      dhcp4: false
      dhcp6: false
      routes:
      - to: default
        via:  WAN_GW
    xconn:
      interfaces: [eno2]
      dhcp4: false
      dhcp6: false
    mgmt:
      interfaces:[eno3]
      dhcp4: false
      dhcp6: false
    back:
      interfaces: []
      dhcp4: false
      dhcp6: false
  version: 2
```

Important

The following sections ("Creating the vNETcs" and "Creating the SDLCs") should be used as a reference when running the standard install procedure. It is not intended to change the order of the standard install process, just to provide guidance to information used during installation.

CREATING THE VNETCS

Follow the Verity quickstart instructions to install the vNetC under KVM.

In the vNetC xml:

- change the name to 'vnetc1' or some appropriate name that indicates this is node 1
- replace the three `<interface>` definitions with these definitions (newer xml files may already be in this format):

```
<interface type='bridge'>
  <source bridge='wan'/>
  <model type='virtio'/>
  <link state='up'/>
```


Copyright BE Networks 2026

```
</interface>
<interface type='bridge'>
  <source bridge='xconn'/>
  <model type='virtio'/>
  <link state='up'/>
</interface>
<interface type='bridge'>
  <source bridge='mgmt'/>
  <model type='virtio'/>
  <link state='up'/>
</interface>
<interface type='bridge'>
  <source bridge='back'/>
  <model type='virtio'/>
  <link state='up'/>
</interface>
```

Continue to "virsh define" the vnetc and start the vnetc per the standard instructions.

If the management and WAN networks are combined, make the Management Address field is cleared in the Admin Network Settings tile,

otherwise set:

- Management Address: `MGMT_NET`

The same procedure is used to create the second vNetC, with appropriate variable changes (eg `VNETC2_IP` instead of `VNETC1_IP`)

Creating the SDLCs

For all configurations:

- the first interface to type 'bridge', source bridge='back'
- the second interface bridge to source bridge='mgmt'
- the third interface is not used and should be removed

Follow the instructions to install the first SDLC for KVM

- Use admin/wizard to set `SDLC_HOSTNAME`, `SDLC_IP`, `ACS_IP`, NTP address will be `VNETC1_IP`
- Use admin/redundancy/serial to set `SDLC_SN`

The same procedure is used to create the second SDLC, with appropriate variable changes (eg `VNETC2_IP` instead of `VNETC1_IP`)

> **Important**
>
> Do not proceed to the next steps until you have both systems installed and running the latest provide core and firmware loads for Verity Release 6.4. Reminder that the SDLCs need to remain isolated from each other.

CONFIGURE VNETCS FOR HA MODE

For both vnETCs, on the console, edit `/etc/rc.conf` with the appropriate vNETC specific IP addresses:

- Change the **vtnet1** IP to `VNETC1_XCONN_IP`
- Add an entry: `ifconfig_vtnet3="inet VNETC1_BACK_IP "`

Create the file "/var/ns/arg_defaults/ns_httpd.conf" and add:

```
append_allowed = 'FQDN'
```

If the GUI will be accessed via an IP address rather than an FQDN, use:

```
append_allowed = 'FQDN,FLOATING_IP'
```

Replicate the host keys from vNetC1 to vNetC2, eg from vNetC1:

Copyright BE Networks 2026

`rsync -ai /etc/ssh/ssh_host_* VNETC2_IP:/etc/ssh/`

Create an ssh key for connections between vNetCs, on vNetC1:

```
ssh-keygen -t ed25519
```

(Hit enter to bypass default location and passphrase)

Copy the generated key to the other vNetC:

On vNetC1:

```
rsync -ai ~/.ssh/id_ed25519* VNETC2_IP:.ssh/
```

This creates a common key for the two nodes to communicate with each other. This is required to allow for synchronization of configuration and statistics.

On vNetC 1, create `/var/ns/ha/failover.conf` with:

```
:include "/usr/ns/etc/failover.conf"
enabled = 1
floating_address = FLOATING_IP
xconnect_address = VNETC1_XCONN_IP
peer_xconnect_address = VNETC2_XCONN_IP
peer_mgmt_address = VNETC2_IP
verification_address = VERIFICATION_IP
sdlc_mgmt_address = SDLC1_BACK_IP
```

On vNetC 2, create `/var/ns/ha/failover.conf` with:

```
:include "/usr/ns/etc/failover.conf"
enabled = 1
floating_address = FLOATING_IP
xconnect_address = VNETC2_XCONN_IP
peer_xconnect_address = VNETC1_XCONN_IP
peer_mgmt_address = VNETC1_IP
verification_address = VERIFICATION_IP
sdlc_mgmt_address = SDLC2_BACK_IP
```

Other variables that alter the default behavior are as follows. Where values are shown, these are already used as the defaults and the variables do not need to be included in /var/ns/ha/failover.conf.

Settle Timing

This specifies the period in seconds before state change actions are taken after decisions have been reached.

A rule-of-thumb is to make the settle period be the reboot time for the server hardware plus 60 seconds.

```
settle = 300
```

Probe Cycle

This is the cycle period for probing the peer. Probing is low-overhead so can be quite short, but should be more than 3 seconds

```
probe_cycle = 10
```

Replication Period

How frequently in seconds to run replication of low priority data from the active to the inactive node. This data include device statistics, non-database configurations like NTP and resolver settings, and forensic application data for ns_bizd, etc.

```
replication_period = 1800
```

Replication Delay

Delay after a failover before replication will be performed.

Copyright BE Networks 2026

```
replication_delay = 300
```

Manual Override

Manual_override should be empty for normal operations. If true, this system will become active, if false this system will become inactive. If this setting is made on just one node, the other node will transition to the appropriate state, assuming communications with this node are operational. (This variable is also administered via ns_admin menu)

```
manual_override =
```

Slave Autoconversion

If slave_auto_conversion is true, a node will be automatically converted to a postgres slave once both the run state and postgres state have been stable for slave_transition_delay seconds.

```
slave_auto_conversion = true
```

Slave Transition Delay

Slave_transition_delay sets how long to wait after a node becomes inactive before attempts a slave conversion.

```
slave_transition_delay = 360
```

Reboot both vNETCs

CONFIGURE SDLCS FOR HA MODE

Use these steps to configure both of the SDLCs via the admin CLI:

- Use admin/network/backdoor to set `SDLC1_BACK_IP` mask 255.255.255.x
- Use admin/security/keys/get - to set an authorized key, taken from /var/ns/ha/sdlc_ha_key.pub in the vNETC

```
 Main/ administration/ Security Base Menu# keys

 Main/ administration/ Security Base Menu/ SSH Keys Base Menu# get -

Enter file (Press enter and <CTRL-D> when done):
ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAAIHHrFWYS5snSDRNV6mN6rZkWMn1Xp3CIXgn4x5HvSJyP root@system.domain.com


Authorized keys 'authorized_keys' copied to '/home/admin/.ssh/authorized_keys'

 Main/ administration/ Security Base Menu/ SSH Keys Base Menu#
```

- Use admin/redundancy/enable to enable redundancy mode

Reboot the SDLC. When it reboots it will be in Standby mode. If this is the first SDLC, shut it down after the reboot until the second SDLC is also in HA mode.

Both SDLCs are set to standby mode waiting for the vNETC to bring online the active SDLC.

As a final check, log into both vNETs and run `ns_info` to ensure that database replication is active on both sides.

INSTALL MONITORING VMS

If the Monitoring VM is used, it is installed first on vNETC 1, pointing to the floating IP address, and then on vNETC2, also pointing to the floating IP address.

Upgrading an HA site

HA sites must be upgraded in a specific sequence, which should be completely without long delays between steps.

This is because during the upgrade, schema changes to the Postgres database may be involved, and these are immediately replicated to the inactive node, which means the BE software on the inactive node may be inconsistent with the Postgres schema.

In addition, it is not possible to do an offline upgrade to the inactive node because upgrades require write-access to Postgres.

Copyright BE Networks 2026

UPGRADE STEPS

- Log into the vNetC console of the inactive node and run ns_admin and select "High Availability Configuration". Select "Force Node Inactive".
- Upgrade the active node via **Admin** -> **Software Packages**.
- After the upgrade has completed, confirm the system has been left in read-only mode. The topology page should display a tan pin-wheel icon. **Leave the system in read-only mode**.
- Use ns_admin on the inactive node to reenable failover. Select "Clear Node Overide".
- Log into the vNetC console of the active node and run `ns_admin`
- Choose "High Availability Configuration" and "Replicate Config and Statistics"
- When that completes, choose "Force node INACTIVE". This will trigger a failover to the currently inactive node.
- Once the GUI is accessible on this node, perform the upgrade in the normal way via **Admin** -> **Software Packages**. If necessary, delete the current package, re-upload it, and re-deploy it.
- Once the upgrade has completed, back on the console of the now-inactive node, use `ns_admin` to:
- Choose "Set Postgres to Slave Mode"
- Select "Clear Node Override".
- After verifying any configuration changes, the system can be returned to read-write mode

Recovering an HA site from POSTGRES backup

When recovering an HA site from a backup, there are extra considerations:

1. PostgreSQL replication must be manually restarted after the restore
2. If the backup came from the other node, the license restored will be for that node, so the correct license must be reapplied.

RESTORE STEPS ON THE ACTIVE VNETC

```
service ns_init stop
ns_restore -b /tmp/yyyymmddThhmmss.dddZ.tbz
ns_ha_postgres -et master
service ns_init start
```

Note that if the backup file was copied from another system, add "-s config" to the ns_restore command to avoid loading license and network configuration from that system, ie:

```
ns_restore -b /tmp/yyyymmddThhmmss.dddZ.tbz -s config
```

RESTORE STEPS ON THE INACTIVE VNETC

To re-enable POSTGRES replication

```
ns_ha_postgres -et slave
```

This re-enables PostgeSQL replication.

How to Verify High Availability Mode of a Deployed Verity Instance

To check the High Availabillity status on a deployed instance of Verity, follow these steps:

1. On the **Main Navigation Bar**, click **Admin**.
2. Click **VNFs**.
3. Double-click the **vNetc Commander** section.
4. Double-click **System Information**.
5. View the **HA (High-Availability) Mode** value. It will be **True** or **False**. 

Copyright BE Networks 2026

4.8.5 Security Lockdown Overview

Overview

This document provides instructions to prepare virtual machines and devices within the Verity platform for operations in highly secure environments. Standards are defined by the Defense Information Systems Agency (DISA). The process involves securing all the components to ensure they are trusted on the network, prevent any possible tampering within the components and that all communications between components use highly secured encryption technologies. All physical devices and the SDLC application are processed independently in the first steps and **should remain disconnected from the system** when the last step, the vNETC orchestration platform is processed and locked down.

4.8.6 Security Lockdown for the vNetC Orchestration Platform

The next step in the process requires the vNetC virtual machine to be installed under VMware.

During the installation you install the following certificates: - Web Certificate - Radius Client Certificate - Certificate Authority Chains

Information

Locked-down vNETCs provide FIPS-140-2 conforming encryption on all networking connections traversing virtual machine and hardware boundaries. All participating devices must provide mutual authentication.

There are other considerations to meet the FIPS-140-2 standard, including the use of the VMware ESXi hypervisor on specific certified hardware. This document assumes that the system is being configured according to the document: **BEVN Installation for VMWare ESXi**.

BEVN Installation for VMWare ESXi

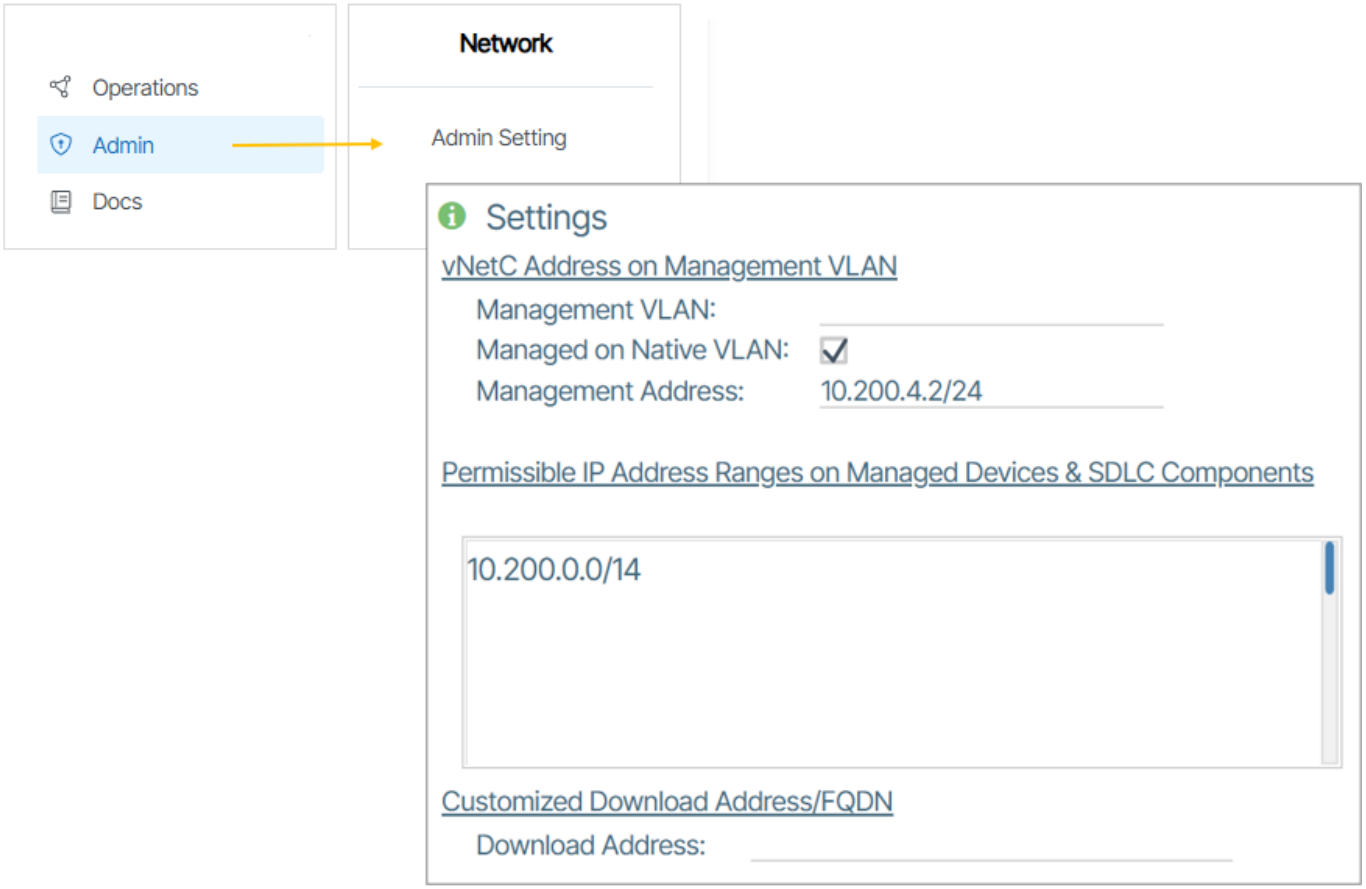
Read VMWare document up to and not including step 7. Instructions here are to be elaborated on and this is placeholder text.

Set Download Address to vNetC FQDN

In secure lock-down, manually locked-down devices must use FQDN to access the vNetC for downloads. Using an IP address is not possible due to TLS certificate requirements.

Go to **Admin** and click the **Admin Settings** option. Set the customized **Download Address** to the vNetC FQDN.

Copyright BE Networks 2026



Certificate Management

vNetC Certificate files are installed from SD-ADMIN via the Certificate Management collection of certificate panels .

Information

For more information see Certificates

VNETC WEB-SERVER CERTIFICATES

The vNetC requires a web server certificate. This certificate provides TLS validation ensuring secure communication when web browsers and devices contact the vNetC. The certificate must fit these criteria.

1. The certificate contains the private key, has FQDN set as the common-name, and a certificate authority chain.
2. The certificate is uploaded to the **vNetC Server Certificate** panel at **Admin/Certificates**.

CERTIFICATE AUTHORITY CHAINS

The vNetC must be loaded with certificate authority chains composed of intermediate and root certificates. For systems in secure lockdown, all **Certificate Chains** are used to validate client certificates (such as CAC) and device client certificates. All such certificate chains should be uploaded to **Admin/Certificate Chain** panel as shown here .. Multiple files can be uploaded and are used when validating all TLS client connections (users or devices)

RADIUS CLIENT CERTIFICATE

The Radius client certificate authenticates the vNetC with the RADIUS authentication server(s). The RADIUS authentication server is configured under **Admin/User Radius Logins** panel as shown in the following diagram .

Copyright BE Networks 2026

The vNetC must have a client certificate file containing the private key and a certificate signed by an authority recognized by the RADIUS server. The certificate is required for systems in secure lockdown, and until it is loaded, only the emergency admin user will have access to the system. For other systems, the certificate is used if provided.

The RADIUS client certificate is uploaded to the **Admin/User Radius Logins** tile  .

REVOCATION LIST

The **Admin/Certificates/Certificate Revocation Management** contains a list of "X509 CRL" entries that revoke client certificates that are no longer valid for use by web browsers or devices  .

vNetC Secure Lockdown Mode

Once the system is configured with the required certificates you perform the following actions:

1. Login as Root user with root password.
2. Run the following command: **ns_admin**
3. Select **SD-LAN Features**
4. Enable **Security Lockdown**.
5. Save, exit and reboot.
6. When the vNetC becomes accessible, you are required to provide a PIN for CAC based authentication before being able to access the login page. Complete the PIN submission action.
7. If you are then able to access the login page and login in, the vNetC is secure and can be made accessible from the production network.

When in Lockdown Mode

- A client certificate from an approved authority is required.
- The emergency username is always "admin", and its password is set during system installation.
- If no RADIUS servers are enabled, or none can be reached, only the emergency user will have access.
- If **any** RADIUS servers are accessible, the emergency user will not have access **unless** the RADIUS server gives authorization.
- All other users must be authorized via the RADIUS server(s).
- The three-login failure limit leads to a 15 min cooling-off period.

Auto-Onboarding (Onboarding Mode)

Verity now has a new feature named **Onboarding Mode**.

In locked-down systems, devices within the appropriate IP address range will not be recognized as managed devices and will not attempt to connect to the vNetC. However, in systems that are not locked-down, the system will recognize such devices as managed devices. In a locked-down system, enabling **Onboarding Mode** creates an exception, temporarily allowing devices with the appropriate settings and IP ranges to be recognized by the system.

Note

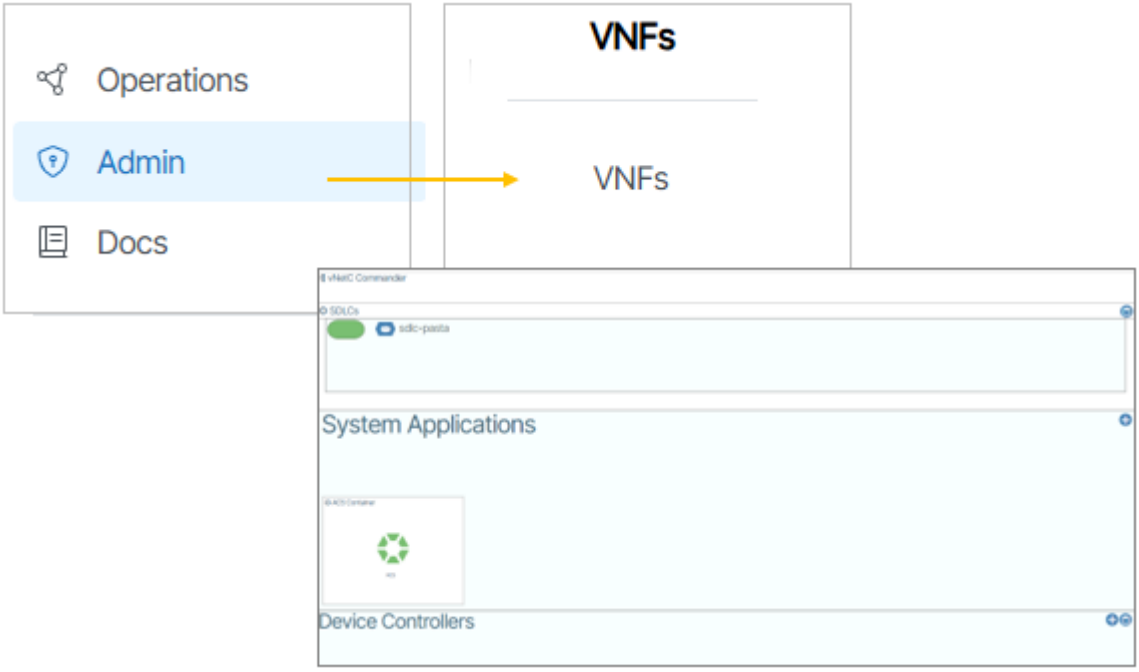
- Device onboarding mode allows "haproxy" on the vNetC and ACS to accept requests from devices without a client certificate.
- Devices with valid certificates (already locked-down) will not be affected.
- Devices with invalid certificates will be rejected, and the only way to recover them is through a hard factory reset while device onboarding mode is enabled.
- There is no current way to know that a device has been rejected or why, as this is managed by "haproxy". This information might be found in /var/log/haproxy.log.
- Devices without a certificate will be allowed to connect, and vNetC will handle the process of delivering a device-specific client certificate and signing certificate for vNetC and ACS, then place the device in lock-down mode.

Copyright BE Networks 2026

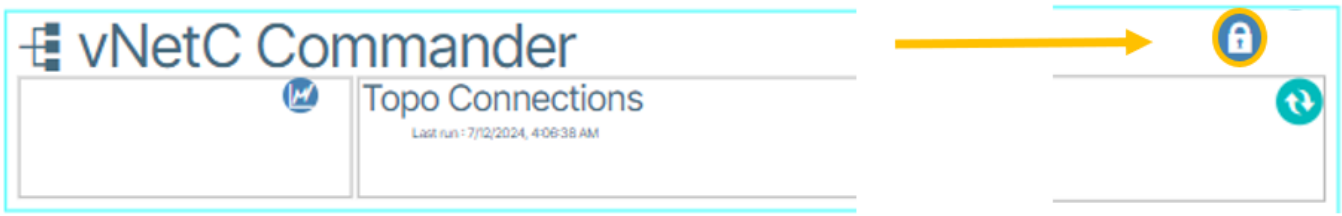
When Onboarding Mode is enabled, each device to be onboarded should be power-cycled to ensure it collects its startup configuration from the vNetC. When the system is not in Onboarding Mode, access to the startup configuration is blocked until a device has all the appropriate certificates.

ENABLE ONBOARDING MODE

To use **Onboarding Mode,** click **Admin** and select **VNFs**.

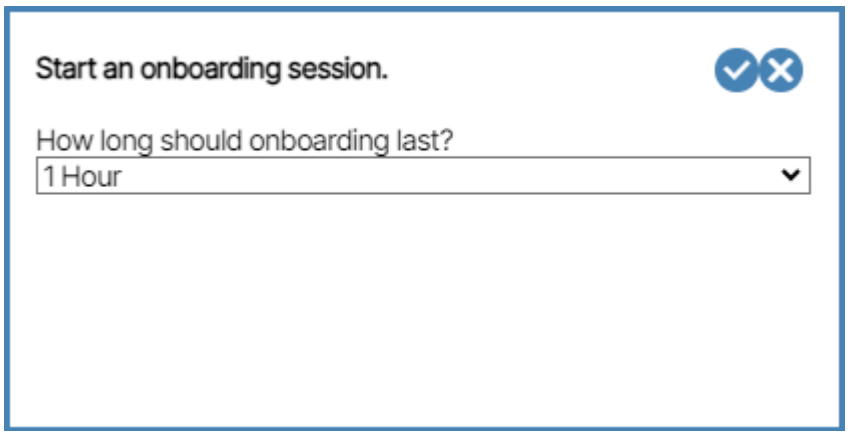


Zoom into the **vNetC Commander** tile and scroll all the way to the right until you see the **Lock** icon .



To enable **Onboarding Mode**, click the Lock icon and place it in the *unlocked* state .

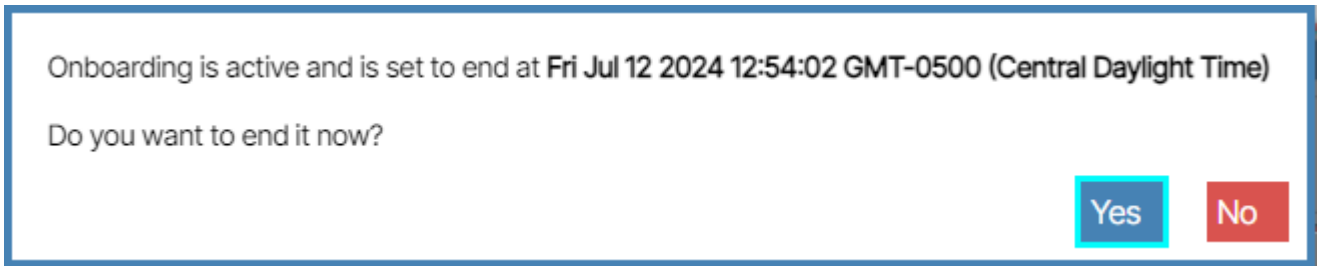
A dialog box appears asking how long the onboarding should last. What you choose determines how long the **Onboarding Mode** is enabled for.



Copyright BE Networks 2026

TERMINATING ONBOARDING MODE

If you click the unlock ( ) icon prior to the **onboarding session** expiration, you are provided with a message asking if you want to preemptively end the session.



Manual Onboarding

To use certificates from an external certificate authority you must manually onboard each device. Detailed instructions are available here.

Campus Lockdown

ADD ACS CERTIFICATE (ADDENDUM)

This is needed when using ONTs FQDN in ACS url.

1. Go to **Admin/Certificates/ACS Server Certificate**
2. Click on **ACS Server Certificate** box.
3. Drag and drop the **server.pem** file.

Copyright BE Networks 2026

4.8.7 Verity Security

This document serves as an overview of the security features integrated into Verity networking software and UI application.

VNetC

The vNetC supports local authentication and remote authentication via RADIUS. In both cases users may be granted these access levels: admin, read-write, and read-only. An additional attribute "tech" can be granted which exposes expert-level options in the user interface.

Passwords

Local Users - Local user passwords are validated against a sha256-based one-way hash. Passwords are managed from the **Administration -> Users -> Accounts** panel or on the external RADIUS server. **System** - Because controllers, switches, and the RADIUS and SNMP agents must be able to recover the clear-text for passwords, it is not possible to use one-way hashes to these store passwords. Instead, they are encrypted using a master key stored on the vNetC using salted AES-256-CBC cryptography with PBKDF2 ("Password-Based Key Derivation Function 2").

Restricting console logic (ssh access requires authorization key)

Restricts access to all components to SSH key only, i.e., no username/password allowed. Enable it from the ns_admin menu.

EXPOSED PORTS

In normal operation the following ports are exposed on the vNetC:

External (WAN) network:

Port	Protocol	Service	Description
3,4,8,11	ICMP	ping	Allows external sources to ping vNetC
22	TCP	SSH	Access allowed from private
68	UDP	DHCP	Receives DHCP answers from broadcast source 67
80	TCP	HTTP	Used for one-time-password image download operations for upgrades. These can optionally be done over HTTPS
161	TCP or UDP	SNMP	External SNMP access
443	TCP	HTTPS	Used for: Web access for users. Secure web socket (wss) user REST-based API access. Device access for discover and status.

Management Network (if configured):

Port	Protocol	Service	Description
123	UDP	ntp	Support NTP sync to devices
1581	TCP	ACS	Used by ACS system to query database
1582	TCP	ACS	Used by ACS system to receive events
1958	TCP	MACs	Receives MAC address table updates
8001	TCP	Stats	Receives statistics and performance metrics

If the vNetC is configured without a separate management network port, these ports are exposed on the WAN interface, restricted to private (RFC1918) addresses.

Copyright BE Networks 2026

RFC1918 addresses are explicitly:

- 10.0.0.0/8
- 172.16.0.0/12
- 192.168.0.0/16
- 169.254.0.0/16

TRANSPORT ENCRYPTION STANDARDS

All encrypted transports (SSH, HTTPS, WSS) meet US FIPS-140-3 security requirements for cryptographic modules.

HTTPS AND SERVER WEB CERTIFICATES

The vNetC HTTPS and WSS access terminates in an "haproxy" HTTPS process. This is configured with these cipher rules for TLSv1.2: EECDH+ECDSA+AESGCM:EECDH+aRSA+AESGCM:EECDH+ECDSA+SHA384:EECDH+ECDSA+SHA256:EECDH+aRSA+SHA384:EECDH+aRSA+SHA256:EECDH:EDH+aRSA:HIGH:!ECDHE-RSA-NULL-SHA:!ECDHE-ECDSA-NULL-SHA:!ADH:!aNULL:!eNULL:!RC4:!DSS:!MD5:!DES:!3DES

Additionally, it utilizes these cipher suites for TLSv1.3: TLS_AES_256_GCM_SHA384:TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256

These ensure a high level of encryption that is compatible with all modern browsers.

SSLv3, TLSv1.0 and TLSv1.1 are disabled.

It is recommended that vNetC's have FQDN and that administrators install a server certificate for the FQDN. When a valid certificate is provided. Strict-Transport-Security is active making SSL-stripping attacks improbable.

We test our TLSv1.3 configuration on an ongoing basis with a vNetC deployed on the public internet, using the Qualys SSL Labs system at SSL Server Test (Powered by Qualys SSL Labs) . You can check our system using the hostname "vnc-sslcheck.beyondedgenetworks.com". Please let us know if this does not return an A+ rating.

SDLC

EXPOSED PORTS

In normal operation the following ports are exposed on the SDLC:

Port	Protocol	Service	Description
830	TCP	Netconf	Allows provisioning and monitoring by the ACS and Vnetc
443	TCP	HTTPS	Allows limited debugging via webpage
22	TCP	SSH	Allows full debugging via CLI
80	TCP	HTTP	Not used
123	UDP	NTP	Used to sync time
67	UDP	DHCP	Used to assign IP addresses if DHCP server is enabled.

Copyright BE Networks 2026

ACS

EXPOSED PORTS

In normal operation the following ports are exposed on the ACS:

Port	Protocol	Service	Description
5555	TCP	TR69/HTTPS	Used for provisioning and monitoring ONTs over HTTPS
5554	TCP	TR69/HTTP	Used for provisioning and monitoring ONTs over HTTP
830	TCP	netconf	Allows provisioning and monitoring by the ACS and Vnetc
514	TCP/UDP	syslog	Used to receive system messages from devices.
8888	TCP/UDP	Vnetc	Used by Vnetc to configure and query the ACS
443	TCP	HTTPS	Allows limited debugging via cli
22	TCP	ssh	Allows full debugging via cli
80	TCP	HTTP	Not used
123	UDP	NTP	Use to sync time
67	UDP	DHCP	Used to assign IP addresses if DHCP server is enabled.

Backend Service(s)

EXPOSED PORTS

In normal operation the following ports are exposed on Backend Service(s):

Port	Protocol	Service	Description
830	TCP	Netconf	Allows provisioning and monitoring by the ACS and Vnetc
514	TCP/UDP	syslog	Used to receive system messages from devices
8080	TCP/UDP	Vnetc	Used by Vnetc to configure and query Backend Service(s)
443	TCP	HTTP	Allows limited debugging via webpage
22	TCP	ssh	Allows full debugging via cli

Device Controller

EXPOSED PORTS

In normal operation the following ports are exposed on the device controller:

Port	Protocol	Service	Description
830	TCP	Netconf	Allows provisioning and monitoring of the by ACS and Vnetc
514	TCP/UDP	syslog	Used to receive system messages from devices
443	TCP	HTTPS	Allows limited debugging via webpage
22	TCP	ssh	Allows full debugging via cli
69	UDP	TFTP	Allows for transfer of configuration for some switches then SNMP protocol is selected for monitoring the switch

Copyright BE Networks 2026

Client Certificate Authentication

The system can be configured to support client certificate authentication. This operates in a "lockdown" mode that meets US DoD requirements. It does impose strict access requirements on all transports in the systems and demands client certificates for all users, devices, and external application access. In "lockdown" mode, TLSv1.2 is disabled and the available cipher rules are reduced to: TLS_AES_256_GCM_SHA384:TLS_AES_128_GCM_SHA256.

System Logs

The system performs extensive logging of internal and device operations. Logging can be viewed from the Alarms history interface.

Serviceability

The system allows download of the current system state across all management components from the web-based user interface. Nonprivileged users can collect this information in an opaque form that can only be accessed by BE engineers. Privileged users can examine the content without BE involvement.

Downloading a System Snapshot

Downloading system state is performed by clicking the Snapshot Button from the Verity Web Interface. Clicking this button reveals a window prompting you to choose the data you want to capture. After you make your selection, you click **Create File** and then **Download**. 

Admin Control Panel

Admin is the user control panel for the VNetC. The **Admin** panel contains tools useful for securing your system.

CERTIFICATES

Certificates is a drag and drop control panel to import SSL Certificates.

> **Information**
>
> For more information see Certificates

SETTINGS

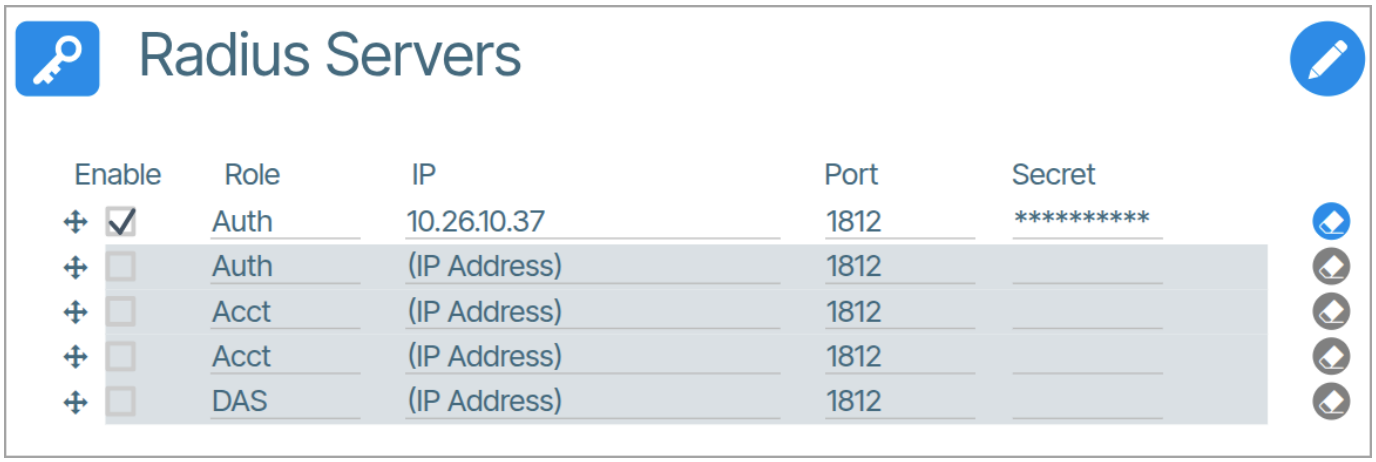
Settings are available in Admin/Network/Settings. Settings lets you configure the following: - vNetc Address on Management VLAN - Permissible IP Address Ranges on Managed Devices - Customized Download Address/FQDN

EXTERNAL APIS

External APIs (Admin/External APIs) lets you configure status monitoring, backups, and loggings of the vNetc via SNMP, SFTP and SYSLOG.

RADIUS Servers

Enable, disable, and configure RADIUS server settings here.



Copyright BE Networks 2026

ALARMS

Alarms are notifications that inform the administrator about current changes to the network state. Active alarm are listed in **Active Alarms**.

Copyright BE Networks 2026

4.8.8 System Recovery

System Recovery from POSTGRES Backup

VERITY 6.3.0 +

If the vNETC is part of a high availability deployment, please refer to the instructions in the High Availability configuration section.

Recovering an HA site from POSTGRES backup

Otherwise, proceed as follows:

1. Transfer a database backup of the system being recovered to the new system. This would come from the SFTP backup server if configured, or from the old system under the directory: `/var/ns/backups/postgres/active/archive/` . The file will typically have a name like: **20240514T164709.116Z.tbz**
2. Place the backup file in a convenient location, such as `/tmp or /var/ns/direct` .
3. Shut down the operations services using: `service ns_init stop`
4. Run the restore command specifying the transferred file, eg: `ns_restore -b /tmp/20240514T164709.116Z.tbz` . This restores the database and system configuration from the database backup.
5. Once complete, restart system services using: `service ns_init start`
6. `virsh console` into SDLC and run the setup to point to the vNetC.

Copyright BE Networks 2026

4.8.9 Debugging

Snapshot files

For most problems, system administrators can use snapshot files to diagnose and correct software issues.Snapshot files exist for the vNETc, UI and all VNFs (Virtual Network Functions). For all reported problems, it is recommended that the vNETC snapshot and the snapshot from the affected component are included in the ticket. The vNETC snapshot contains sections for both the internal processes as well as a UI snapshot. The system gives you the option of various size files. The default is recommended unless you are trying to capture an event that may have occurred several days in the past.

WHAT INFORMATION IS STORED IN SNAPSHOTS?

1. Device Controllers - the Device controller snapshots are retrievable from the top section of the switches in the Topology View of the UI, or within the Device Controller object within the VNF section. For most provisioning or device related issues, it is recommended that the ticket includes the vNETC snapshot as well as the snapshot of the affected device controller(s).
2. Backend Services - provides a memory resident UI client for caching status and reporting information. In the event that the UI is displaying erroneous status, Backend Services snapshot should be provided along with the vNETC.
3. SDLC - the SDLC is responsible for creation and monitoring of the VNFs inside of Verity. The snapshot contains information regarding the provisioning of the VNFs from the VNETC and the processes that create and monitor the VNFs.
4. ACS - the ACS is the primary communications interface between the vNETC and all devices controllers. Any simultaneous communications issues between the vNETC and devices controllers and other VNFs would necessitate including an ACS snapshot.

CREATING SNAPSHOTS

1. Click **Support** from the left main navigation menu.
2. Click **Snapshot**.
3. In the window that appears, click the attributes of the application that you want to include in the snapshot. There are options for VNC, Observability VM, AI Ops VM, and GUI.
4. Describe your issue by typing a summary of the problem in the text box.
5. Click **Create File**. The file will be created on the application server.
6. Click **Download File**. This file is downloaded to your computer and can be sent to customer service.

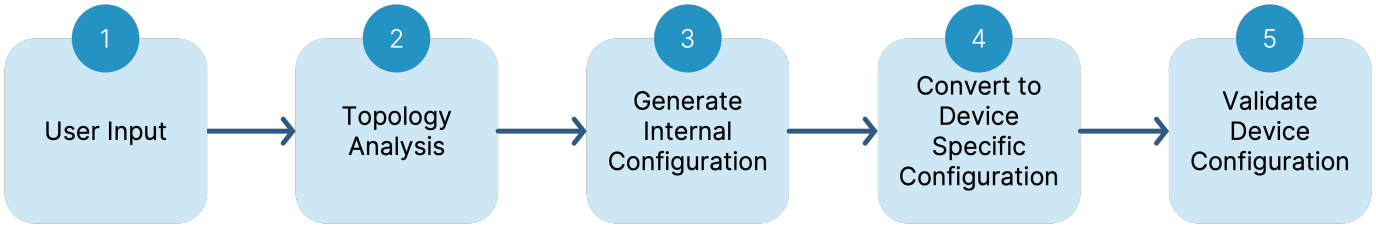
Provisioning Problems

The main function of Verity orchestration is to translate the user's intent into switch configurations. This is done through several stages of data model and protocol translations. In the event that the intended configuration does not end up properly on the physical device or if the device shows a provisioning failure (entire switch is red) there are various tools available to debug the issue through the Verity components.

To summarize the process, Verity software consumes the users's intent from the managed objects exposed to the user in the UI and on the Northbound Interface APIs. It combines this information with network topology retrieved from the managed devices in order to generate device configurations that are passed between the vNETC and the Device Controllers. The Device Controllers, in turn, translate Verity's internal data model to the protocol supported by the switch. Examples of these protocols include SNMP, gNMI, NETCONF as well as simple CLI commands in some cases.

Copyright BE Networks 2026

Provisioning Flow



1. User Input -->
2. Topology Analysis -->
3. Generate Internal Configuration -->
4. Convert to Device Specific Configuration -->
5. Validate Device Configuration -->

Verity automatically detects events related to user input and topology changes and all actions in the flow should be automatic. Events are passed between processes to make this automation seamless. There are cases during heavy management traffic floods, or in the case of software bugs that the automatic update events are lost. The user can view the current configuration state of the various steps in the flow as well as interrogate the managed switch directly.

THE CURRENT CONFIGURATION OF THE SYSTEM IS VIEWABLE AT THE VARIOUS STAGES OF THE PROVISIONING FLOW

1. UI displays managed objects as provisioned by the user in the various sections of provisioning templates
2. The switch has a drop down menu on the map that you can click to see the various configurations

a. "Running Config" shows the internal Verity Model as built by Bizd and sent to the switch controller

b. "Target running" shows the last uploaded config from the physical device to the device controller

c. "Target Pending" shows what changes are planned when the switch is removed from read only mode

3. The physical switch is accessed directly and shows its running configuration

Manual Resynchronization

THERE ARE FIVE MANUAL RESYNCHRONIZATION POINTS AVAILABLE FOR THE PROVISIONING FLOW IN THE EVENT THAT THE CONFIGURATIONS AT THE VARIOUS STAGES APPEAR TO BE OUT OF SYNCHRONIZATION.

1. Two related to topology calculations (step 2 above) referred to as "Topo Connections" and "Topo Switch Trees" in the vNETC Commander section of the VNF Pane
2. Internal configuration calculation referred to as "Bizd" in the same VNF section
3. Force vNETC to Device Controller Resync (circular arrow on the Switch in the UI network view)
4. Force the Device Controller to resynch with the physical switch (from the current switch configuration button on the switch, there is a circular arrow as well)

Generally speaking, if any of the manual resynchronization attempts cause an update or correction to the subsequent steps in the flow, an issue should be raised to customer support. Snapshots for the vNetC and the affected switch/device controller should be provided.

Copyright BE Networks 2026

5. API

API REST

5.1 Verity API

Verity API is a REST API conforming to the OpenAPI 3.0 specification used for operating a Verity instance.

The Verity API provides developers with access to essential functionality for interacting with data and services within the Verity system. Users can leverage this API to authenticate users, retrieve user information, and manage overlay resources efficiently. With the Verity API, developers can seamlessly integrate Verity's robust features into their applications, enhancing user experiences and streamlining workflows.

API Documentation Versioning

The following documentation provides a high level explanation of how the Verity API functions. For detailed examples and API reference for your exact version of Verity, refer to the Swagger documentation page on your Verity server.

5.1.1 API Documentation

1. Navigate to `/swagger` on your Verity server IP. **Example**: `https://vnc-datacenter.company.net/swagger`
2. On this page, in the top left corner, is a link to the schema file if you wish to download it. To download a JSON copy of the API specification right-click on **swagger.json** and choose Save-As. 

5.1.2 Quickstart

There are **two endpoints** you need to understand before using the API in production:

1. Authorization Endpoint

The **Authorization** section of the Verity API allows users to send their user/password combination to Verity for the creation of a time-limited access token. Once this token is retrieved, all further API access needs to use the token.

CREATE AN AUTH TOKEN

1. Go to **Authorization** and click the drop down menu. 
2. Click **Try it out**. 
3. In the request body, type the **username** and **password** credentials replacing the default string text. 
4. Click the **Execute** button to invoke the HTTP request. 
5. If successful, the action returns a 200 status code with a description of **Successful authorization**. The token is required for requests and expires in 30 minutes. The **Authorization** drop down menu does not need to remain open and can be closed.

Copyright BE Networks 2026

2. Changesets Endpoint

CREATE A CHANGESET

To use a changeset, users must first create the changeset via `POST/changesets`.

1. Go to `POST/changesets`.
2. Unfold the field and click **Try it out**.
3. To *create a new changeset*, the **commit** must be set to **false**, and the **changeset_name** must be set to a unique name. 

> **Changeset Naming**
>
> Creating a changeset with the same name as a pre-existing changeset will return an error.

COMMIT A CHANGESET

1. To commit a changeset to the system, the **commit** field must be set to **true**, and the **changeset_name** must be set to the changeset the user wants to commit. 
2. Click the **Execute** button.

VERITY CHANGESETS (UI)

Changesets are listed within the UI at **Operations -> Change Sets**. 

CHANGESET ACTIONS (UI)

- **Compare to Current**: Compares the changes in the changeset to the current system state in the form of a tabular summary.  Clicking **View Details** provides more detailed information in the form of a **Report**.
- **Download**: Downloads a JSON file of the affected object(s) to the user's machine.

5.1.3 External Integrations

Customers are free to develop their own integrations and plugins that make use of the Verity API. A number of prebuilt options are available:

Copyright BE Networks 2026

API REST

5.2 API Definition

Copyright BE Networks 2026

6. Reference

6.1 How to Use This Documentation

When reading these pages, please note the following common references:

6.1.1 Icon Links

There are many different UI icons, and when following steps in these docs you may see the icon that is being referenced in this notation:



6.1.2 Screenshot Links

When completing a step in a task it is sometimes hard to locate the area of the UI being expressed. Example screenshots of UI locations are noted as:



Clicking these icons results in a full size image.

6.1.3 Error Reporting

If you encounter any problems with this documentation please contact customer service with suggestions.

Copyright BE Networks 2026

6.2 Terminology

6.2.1 Verity Managed Objects

Term	Description
vNETC	Virtual Network Orchestration platform.
SDLC	Application that hosts containers associated to switches and other system functions.
Device Controller	A container within the SDLC application used to manage a switch.

6.2.2 General Networking Terms

Term	Description
Spine-Leaf	A data center network topology that consists of two switching layers— a spine and leaf. The leaf layer consists of access switches that aggregate traffic from servers and connect directly into the spine or network core. Spine switches interconnect all leaf switches in a full-mesh topology.
Pod	A set of connected spines and leafs.
Superspine	A third switch layer that is used to connect Pods to each other.
Underlay	The layer 3 VxLAN based connections between the spines and leafs in the data center.
Tenant	A logical separation of services within the datacenter architecture. The tenants' connections constitute the **overlay** of the datacenter architecture.
Management Switch	A switching device used to provide connectivity between the Verity applications and the out of band management connections of the spines and leafs.

Copyright BE Networks 2026

6.2.3 Functionality Terms

Term	Description
Service	A layer 2 Ethernet model of VLAN identification, QoS, multicast and other layer 2 related options.
Endpoints	(Managed Endpoints, Service Endpoints) – The “edge” of the network where services connect to client devices. Examples of service endpoints are a room number, street address, office desk location or pole location. Service Endpoints specifically are not a physical network device, but have an “Edge Device”, such as an ONT or Ethernet switch, associated with them.
Edge Device	a physical unit, such as an ONT or an Ethernet switch, that is fully managed by the Verity system. The ports of the Edge Device deliver the services to the client equipment.
Top of Rack Switch (TOR)	The layer 2 device that is at the top of the access network topology. It connects to services and routers above, and to the Aggregation Switches and Edge Devices below. In redundant systems there are two TORs, referred to as TOR A and TOR B.
Aggregation Switch (Agg)	The layer 2 devices that connect the TOR to the Edge Devices. There can be multiple layers of Agg switches. Ports on Agg switches may be used as Service Endpoints as well.
Switch Endpoint	the database object associated with switch fabric devices (TOR and Agg).
Service Port	entry point of a layer 2 service into the access network normally through the TOR switch.
Service Drop	Ethernet port at the edge of the access network that connects to client devices.
Direct Endpoint	a Service drop directly egressing the network from the fabric switch complex (TOR or AGG) without using an Edge Device. Usually for craft PCs, back-office connections, trunks to unmanaged devices outside the scope of Verity.
White Box Switch	commercial, off the shelf switch fabric hardware installed with Network Operating System (NOS). The NOS uses standard ONIE installer methods.
B**lack Box Switch**	any “non-open” switching gear managed by the Verity system. In other words, a TOR, Agg or Edge Device that cannot use Verity NOS. They are usually managed internally by the Verity system using SNMP, CLI or NETCONF.
ONT	Optical Network Terminal (aka Optical Edge Device) – an Edge Device providing connectivity at the Endpoint that has an optical WAN interface and multiple service drops.
SDLC	The controller virtual application that hosts containers associated to BlackBox switches.
Device Controller	A container within the SDLC application used to manage a “black box” switch as a TOR, Agg or Edge Device.

Copyright BE Networks 2026

- Templating Terms – (the flow of these items is shown on the configuration map)
- MAC Filters – set of parameters related to MAC Addresses and Hexadecimal MASKs
- Ethernet Port Settings – set of parameters related to the physical layer of an ethernet interface (e.g. Speed, full/half duplex, LLDP, PoE etc.)
- Ethernet Port Profiles – combination of Ethernet Port Settings and one or more Services and MAC Filters defining the Layer 1 and layer 2 connection for the service drop. (e.g. High-Speed Internet, Video)
- Authenticated Ethernet Port – list of possible Ethernet Port Profiles to be applied to the service drop based on connect client authentication (e.g. 802.1x, MAC authentication)
- Service Port Profile – collection of services applied to a Service Port
- Voice Port Settings – parameters related to VoIP protocols, interfaces, and related features.
- Device Settings – set of parameters associated to a physical device such as Power-over-Ethernet (PoE) and LED behavior.
- Device Voice Settings – parameters related to the VoIP client installed on an Edge Device.
- Endpoint Bundle – collection of Device Settings, Ethernet Port Profiles, Resident Applications, and Voice Port Settings to be applied to a Managed Edge Device
- Domain – group of sites in multi-site Verity system
- Site – physical site at a location. A site may contain multiple separated layer 2 switch topologies.
- Top of LAN – The switch at the top of a LAN instance within a site with multiple LANs.

Copyright BE Networks 2026

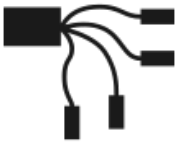
Buttons **Icons**

Copyright BE Networks 2026

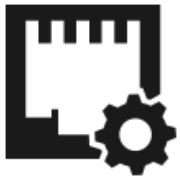
6.3 Display Icons

This page documents all icons available in the icon library.

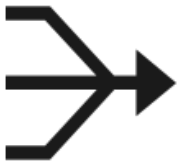
Copyright BE Networks 2026

Icon	Description
	Network Time Traveler Backup
	Migrations
	Badge
	SFP Breakouts
	Endpoint Bundle
	NAC Port Profile, Authenticated Eth-Port
	DHCP Server
	Fabric Collection

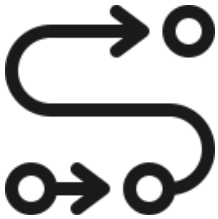
Copyright BE Networks 2026

Icon	Description
	
	Endpoint View, Image Update Sets
	Showtech
	Eth-Port Settings
	SFTP, vNet Commander Network Configuration, Software Package, Service Report, SNMP, SYSLOG, User Account
	Gateway Profile
	Traffic Graph
	Interface

Copyright BE Networks 2026

Icon	Description
	
	Static Connections
	RADIUS Login, Radius
	Lag
	User Type
	MAC Filter
	Subinterface
	ONT, unregisteredObject

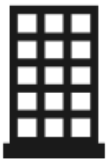
Copyright BE Networks 2026

Icon	Description
	
	Packet Queue
	PB Egress Profile
	Preprovisioned Device
	Preprovisioned Switch
	System Resources
	AS Path Access List, Community List, Extended Community List, IPv4 Filter, IPv4 List Filter, IPv6 Filter, IPv6 List Filter, IPv6 Prefix List, MAC Access List, IPv4 Prefix List, Route Map, Route Map Clause
	Service

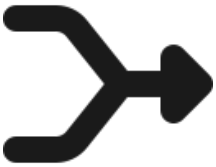
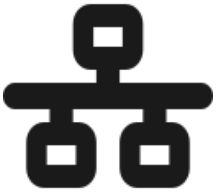
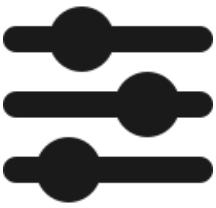
Copyright BE Networks 2026

Icon	Description
	
	Service Port Profile
	Static IP
	Switch
	ACS, AI-Ops
	System Graph
	Device Voice Settings
	Voice-Port Profiles

Copyright BE Networks 2026

Icon	Description
	
	Tenant
	Vnetc
	Gateway
	Tenant
	Eth-Port Profile
	Fabric
	Leaf

Copyright BE Networks 2026

Icon	Description
	
	LAG, OLAG, Olag Channel
	Service
	Device AAA Profile
	Device Settings

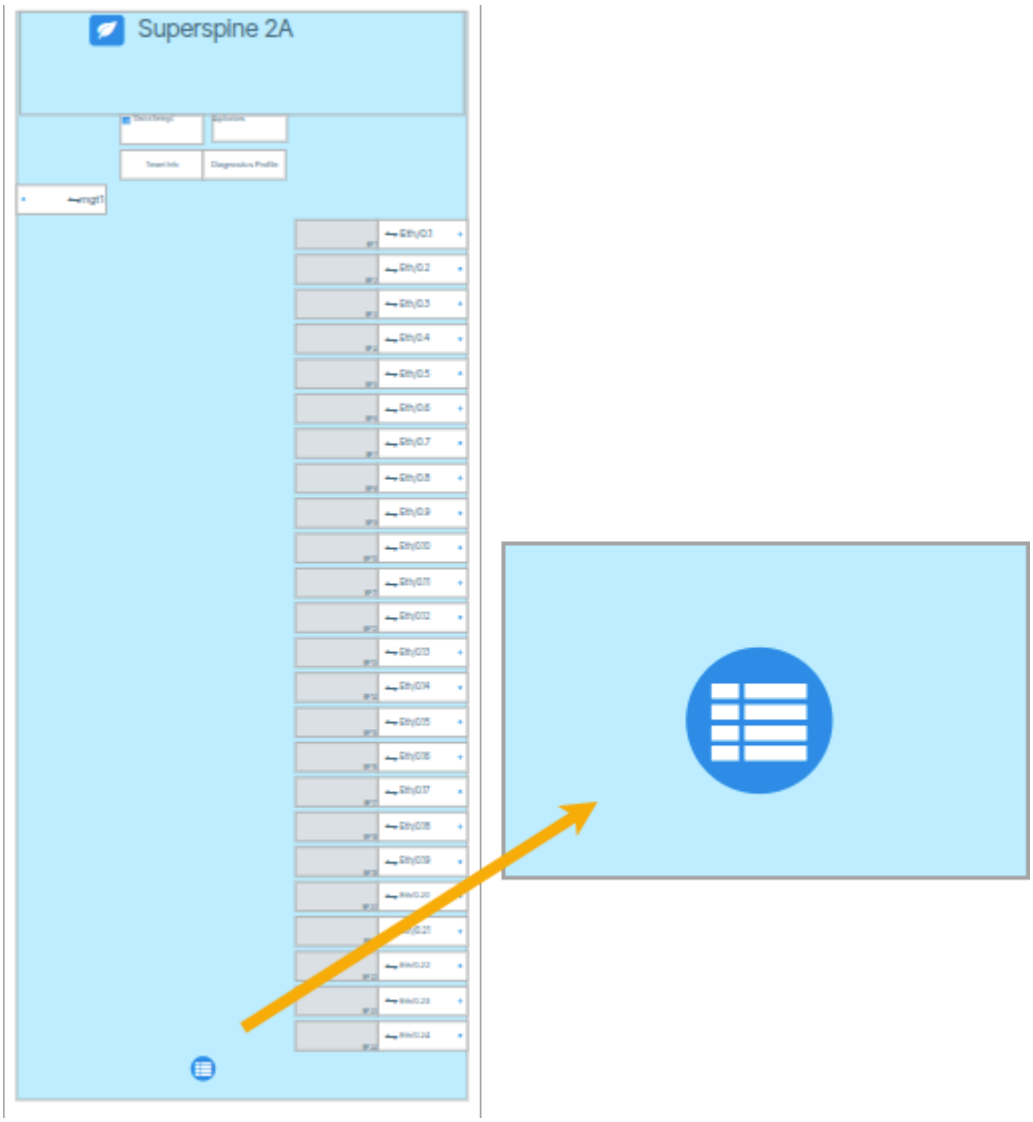
6.3.1 Object Icon Buttons

Object Icon buttons represent tangible objects such as switches or other devices and respond to both right and left click mouse button actions. Left-clicking these buttons causes the window view to center on the button container. Right clicking the button gives you additional options that apply to the object of that button.



Copyright BE Networks 2026

MANAGE ROWS BUTTON



The **Manage Rows** button ( ) provides a convenient way to apply the same settings to multiple ports on a single switch. The settings include switch provisioning as well as the mass enabling/disabling of Edge ports.

Set Eth-Port Provisioning

This option lets you choose an **Eth-Port Profile** and apply it to a port or to a consecutive list of ports.

Set Eth-Port Setting

This option lets you choose an **Eth-Port Setting** and apply it to a port or to a consecutive list of ports.

Replace Eth-Port Provisioning

This option lets you choose all ports by their assigned **Eth-Port Profile** and replace the assignment with a different **Eth-Port Profile**.

Replace Eth-Port Settings

This option lets you choose all ports by their assigned **Eth-Port Setting** and replace the assignment with a different **Eth-Port Setting**.

Copyright BE Networks 2026

Clear All

This setting clears **Eth-Port Provisioning** values and sets **Eth-Port Profiles** to default for a port or a consecutive list of ports. 

Enable Edge Ports

This setting lets you enable an **Edge Port** or enable a consecutive list of **Edge Ports**. 

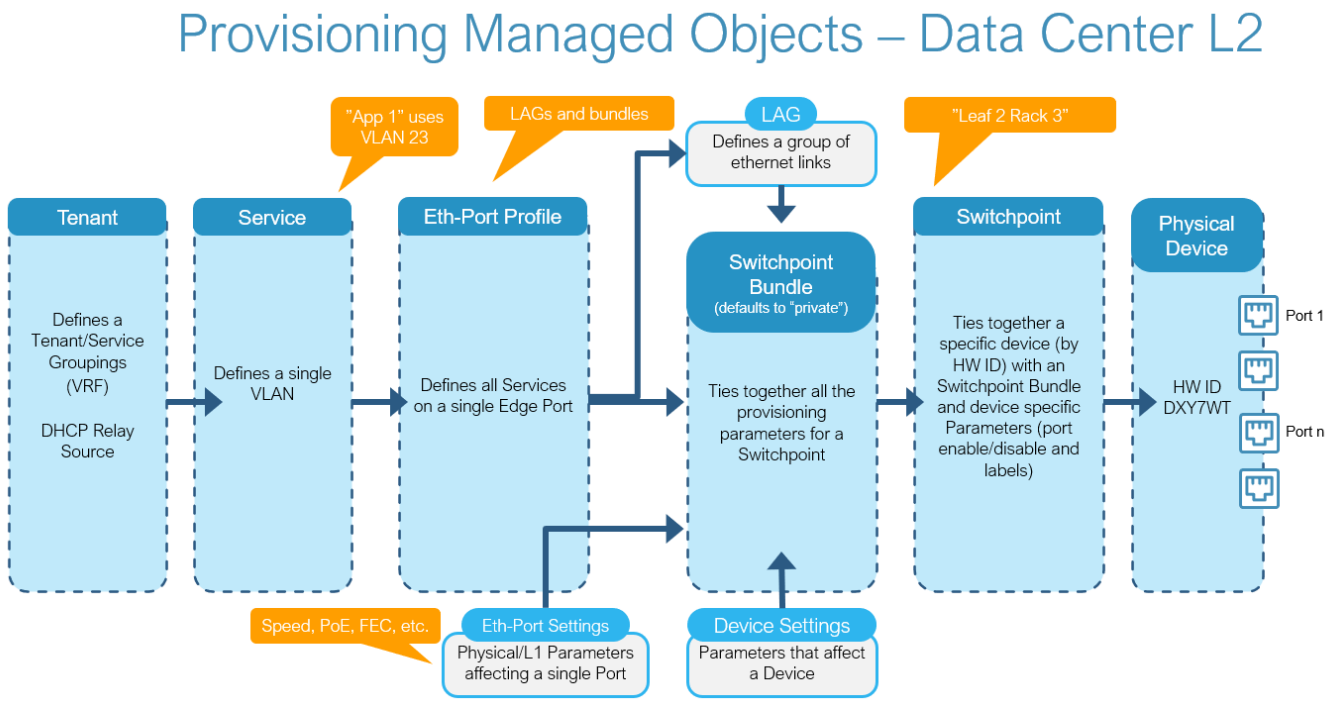
Disable Edge Ports

This setting lets you disable an **Edge Port** or disable a consecutive list of **Edge Ports**. 

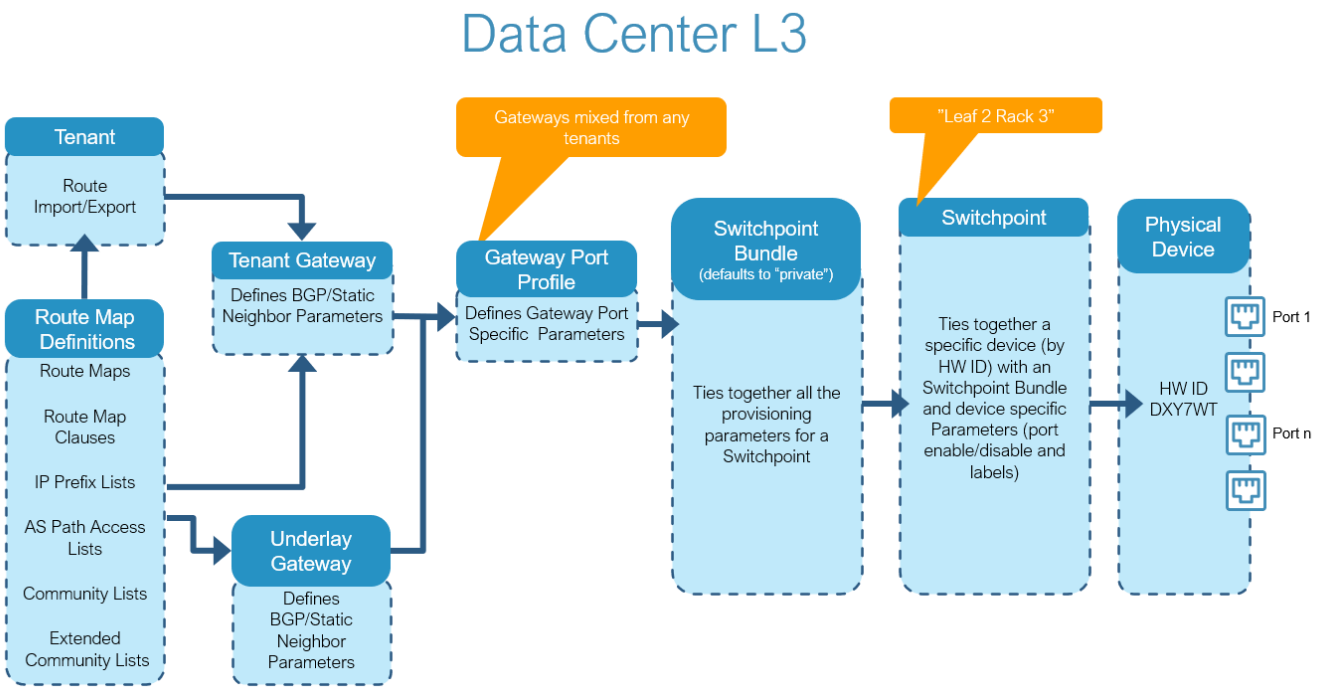
Copyright BE Networks 2026

6.4 Managed Objects

Term	Description
Tenant	A network isolation service implemented as a separate routing domain (VRFs). Tenants contain services and gateways.
Service	A layer 2 VLAN assigned to various ports of leaf switches and interconnected via the underlay.
Ethernet Port Profile	Collection of services to be assigned to Ethernet ports or LAGs
Ethernet Port Settings	Parameters associated with the physical layer of an Ethernet port
LAG	Link Aggregation Group - collection of Ethernet Ports behaving as a single Ethernet interface
Bundle	Collection of the various templates assigned to a Switchpoint
Switchpoint	The abstraction of a managed device within the Verity platform. The database object is responsible for linking switch provisioning to its corresponding physical switch.
Device Settings	Device specific settings
Route Map Definitions	Defintions of rules associated with route manipulations
Tenant Gateway	Template for defining the external connection to a Tenant via BGP or static routes
Underlay Gateway	Template for defining a BGP external connection to the Underlay fabric
Gateway Port Profile	Template for assigning Gateways to Ethernet ports



Copyright BE Networks 2026



Copyright BE Networks 2026

6.5 Downloads / Resources

- Verity VM Sizing Resource Calculator
- FDC Verity Import Example File - Using Switchname as Hostname
- FDC Verity Import Example File - Using Service Tag as Hostname
- FDC Verity Import Example File - Using Static Controller IP Addresses
- FCA Verity Import Example File HGX
- FCA Verity Import Example File Spectrum-X with Dual Planes Servers Unmanaged
- FCA Verity Import Example File Spectrum-X Quad Plane Example
- Alert Manager YML Files (Blank and Example)
- SNMP MIB
- Packet Broker Calculator

Copyright BE Networks 2026

API Parameters

Copyright BE Networks 2026

6.6 Parameters

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
Site	Frames at the bottom of the Network Section	Aggressive Reporting	"true"	Fast Reporting of Switch Communications, Link Up/ Down, and BGP Status
		Anycast MAC Address	"(auto)"	Site Level MAC Address for Anycast
		CRC Failure Threshold	"5"	Threshold in Errors per second that when met will disable the links as part of LAGs
		Enable DHCP Snooping	"false"	Enables the switches to monitor DHCP traffic and collect assigned IP addresses which are then placed in the DHCP assigned IPs report.
		BGP Advertisement Interval	"30"	Leaf BGP Advertisement Interval
		BGP Connect Timer	"120"	BGP Connect Timer
		BGP Hold Down Timer	"180"	Leaf BGP Hold Down Timer
		BGP Keep Alive Timer	"60"	Leaf BGP Keep Alive Timer
		MAC Address Aging Time	"600"	MAC Address Aging Time
		MCLAG Delay Restore Timer	"300"	MCLAG Delay Restore Timer
		BGP Advertisement Interval	"30"	Spine BGP Advertisement Interval
		BGP Connect Timer	"120"	BGP Connect Timer
		BGP Hold Down Timer	"180"	Spine BGP Hold Down Timer
		BGP Keep Alive Timer	"60"	Spine BGP Keep Alive Timer
AS Path	Global Section - Route Map Definitions	Notes	""	User notes describing this provisioning object
		Enable	"false"	Enable flag of this provisioning object
		Enable	"false"	Enable of this AS Path Access List
		Regular Expression	""	Regular Expression to match against BGP Community Strings

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
		Name	""	Name of this provisioning object. Must be unique in its type
		Permit/Deny	"permit"	Action to take on a match of the Community Strings, either Permit or Deny
User Authentication	Admin Tab - Users and Permissions	Overridden Object	"0"	Denotes provisioning object that has been overridden.
		isPredefined	"false"	Denotes this is a predefined provisioning object
		Enable	"true"	Enable flag of this provisioning object
		Name	""	Name of this provisioning object. Must be unique in its type
		Server # Enable	"false"	Enable flag of this Authentication server
		Server # IP	""	IP of this Authentication server
		Server # Port	"1812"	Port of this Authentication server
		Server # Role	"Auth"	Role of this Authentication server
		Server # Shared Secret	""	Shared Secret of this Authentication server
Badges	Global Section - Badges	Notes	""	User notes describing this provisioning object
		Color	"blue"	Color of Badge
		Enable	"true"	Enable flag of this provisioning object
		Name	""	Name of this provisioning object. Must be unique in its type
		Number	"0"	Number of Badge
Device Controller	VNFs Device Controllers	CLI Access Mode	"SSH"	CLI Access Mode
		SNMP Community String	"private"	Comm Credentials
		Comm Type	"snmpv2"	Comm Type
		Co-located VNFs	"routed"	vNetC and SD LAN Controller are both located in the managed compute server

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
		Gateway	""	Gateway
		IP and Mask	""	IP and Mask
		IP Source	"dhcp"	IP Source
		Uses Tagged Packets	"false"	Indicates if the direct interface expects tagged or untagged packets
		Enable Password	""	Enable Password - to enable privileged CLI operations
		Enable	"false"	Enable flag of this provisioning object
		Switch	""	Switchpoint locating the Switch to be controlled
		Upstream is LAG	"false"	If checked, then ZTP will provision the TOR switch with the first 32 ports as a lag to facilitate plug-n-play
		LAG	""	LAG
		LLDP Search String	""	Optional unless Located By is 'LLDP' or Device managed as 'Active SFP'. Must be either the chassis-id or the hostname of the LLDP from the managed device. Used to detect connections between managed devices. If blank, the chassis-id detected by the Device Controller via SNMP/ CLI is used
		Located By	"LLDP"	Controls how the system locates this Device within its LAN
		Switchpoint	""	Switchpoint reference
		Managed on native VLAN	"true"	Managed on native VLAN
		Device managed as	"switch"	Mode
		Name	""	Name of this provisioning object. Must be unique in its type
		Password	""	Password
		Port	""	Port locating the Switch to be controlled
		Power State	"on"	Power state of Switch Controller
		SDLC	""	

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
				SDLC that Device Controller belongs to
		Connection Service	"service.79"	Connect a Service
		Passphrase	""	Passphrase
		Authentication Protocol	"MD5"	Protocol
		Private Password	""	Password
		Private Protocol	"DES"	Protocol
		Security Type	"noAuthNoPriv"	Security level
		SNMPv3 Username	""	Username
		SSH Key or Password	""	SSH Key or Password
		Switch Gateway	""	Gateway of Managed Device
		IP and Mask	""	IP and Mask
		Target Model	"generic_advanced_snmp"	Target Model
		SFP MAC Address or SN	""	SFP MAC Address or SN
		Uplink Port	""	Uplink Port of Managed Device
		Username	""	Username
SFP Breakouts	Global - SFP Breakouts	isPredefined	"false"	Denotes this is a predefined provisioning object
		Enable	"false"	Enable
		Breakout	"1x100G"	Breakout definition; defines number of ports of what speed this port is brokenout to.
		Vendor	""	Vendor
		Part Number	""	Part Number
		Enable	"false"	Enable flag of this provisioning object
		Name	""	Name of this provisioning object. Must be unique in its type
Switch Provisioning	Network View - on each switch	CLI Commands	""	CLI Commands
	Tenants - \<Tenant> - Tenant Slice	Enable	"false"	Enable flag of this provisioning object

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
		Device Settings	"eth_device_profile.76"	Device Settings for this device
		Eth Port # Eth Port Profile	""	Eth Port Profile for this Eth Port
		Eth Port # Eth-Port Settings	""	Choose an Eth Port Settings
		Eth Port # Gateway Profile	""	Gateway Profile for this Eth Port
		Device Voice Settings	""	Device Voice Settings for this device
Community List	Global Section - Route Map Definitions	Notes	""	User notes describing this provisioning object
		Any/All	"any"	BGP does not advertise any or all routes that do not match the Community String
		Enable	"false"	Enable flag of this provisioning object
		Standard/ Expanded	"standard"	Used Community String or Expanded Expression
		Community String/Expanded Expression	""	Community String in standard mode and Expanded Expression in Expanded mode
		Enable	"false"	Enable of this Community List
		Mode	"community"	Mode
		Name	""	Name of this provisioning object. Must be unique in its type
		Permit/Deny	"permit"	Action to take on a match of the Community Strings, either Permit or Deny
DHCP Server	VNFs - System Applications - DHCP Server	Row # Service	""	Connect a Service
DHCP Server		Enable	"false"	Enable flag of this provisioning object
DHCP Server		Excluded Addresses - High	""	Count of excluded addresses that DHCP server will not hand out below the maximum subnet address
DHCP Server		Excluded Addresses - Low	""	Count of excluded addresses that DHCP server will not hand out starting at the base subnet address
DHCP Server			""	Lease duration in minutes

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
		Lease Duration (minutes)		
DHCP Server		Name	""	Name of this provisioning object. Must be unique in its type
DHCP Server		Add acs	""	ACS is a required system component that communicates with ONTs. You choose the SDLC to place it on
Switch Provisioning	Network View - on each switch	Eth # Label	""	Label of this Eth Port
		BGP AS-Number	"(auto)"	BGP Autonomous System Number for the site underlay
		Device Serial Number	""	Device Serial Number
		Disabled Ports	""	Disabled Ports
		Enable	"true"	Enable flag of this provisioning object
		Breakout	""	Breakout
		Switch Router-ID(IP/Mask)	"(auto)"	Switch BGP Router Identifier
		Switch VTEP-ID(IP/Mask)	"(auto)"	Switch VETP Identifier
		Name	""	Name of this provisioning object. Must be unique in its type
		Out of Band Management	"false"	For Switch Endpoints. Denotes a Switch is managed out of band via the management port
		Pod	""	Pod â€“ subgrouping of spine and leaf switches
		Rack	""	Physical Rack location of the Switch
		Read Only Mode	"false"	When Read Only Mode is checked, vNetC will perform all functions except writing database updates to the target hardware
		Type	"leaf"	Type of Switchpoint
Device Template	Provisioning - Device Settings	Enable	"false"	Enable flag of this provisioning object
		Name	""	Name of this provisioning object. Must be unique in its type

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
		External Battery Power Available	"40"	External Battery Power Available
		External Power Available	"75"	External Power Available
		Mode	"IEEE 802.3af"	Mode
		Usage Threshold	"0.99"	Usage Threshold
		Security Audit Interval	"60"	Frequency in minutes of rereading this Switch running configuration and comparing it to expected values.\ if the value is blank, audit will use default switch settings.\ if the value is 0, audit will be turned off.
		Commit to Flash Interval	"60"	Frequency in minutes to write the Switch configuration to flash.\ if the value is blank, commit will use default switch settings.\ if the value is 0, commit will be turned off.
Physical Layer Port Template	Provisioning - Eth-Port Settings	CLI Commands	""	CLI Commands
		Duplex Mode	"Auto"	Duplex Mode
		BPDU Filter	"false"	Drop all Rx and Tx BPDUs
		BPDU Guard	"false"	Block port on BPDU Receive
		Fast Learning Mode	"true"	Enable Immediate Transition to Forwarding
		Guard Loop	"false"	Enable Cisco Guard Loop
		Detect Bridging Loops	"false"	Enable Detection of Bridging Loops
		STP Enable	"false"	Enable Spanning Tree on the port. Note: the Spanning Tree Type (VLAN, Port, MST) is controlled in the Site Settings
		Enable	"false"	Enable flag of this provisioning object
		FEC	"unaltered"	FEC is Forward Error Correction which is error correction on the fiber link.None: Disables FEC on an interface. FC: Enables FEC on supported interfaces. FC stands for fire code. RS: Enables FEC on supported interfaces. RS stands for Reed-Solomon code.

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
		Name	""	Name of this provisioning object. Must be unique in its type
		Allocated Power	"0.0"	Power the PoE system will attempt to allocate on this port
		Enable	"false"	PoE enable
		Priority	"High"	Priority given when assigning power in a limited power situation
Ethernet Port Profile Template	Provisioning - Eth-Port Profiles	Port Monitoring	""	Defines importance of Link Down on this port
		Enable	"false"	Enable flag of this provisioning object
		Name	""	Name of this provisioning object. Must be unique in its type
		Row # Enable	"false"	Enable flag for this row
		Row # Service	""	Choose a Service to connect
		Row # External VLAN	""	Choose an external vlan
		Trusted Port	"false"	Trusted Ports do not participate in IP Source Guard, Dynamic ARP Inspection, nor DHCP Snooping, meaning all packets are forwarded without any checks.
Extended Community List	Global Section - Route Map Definitions	Notes	""	User notes describing this provisioning object
		Any/All	"any"	BGP does not advertise any or all routes that do not match the Community String
		Enable	"false"	Enable flag of this provisioning object
		Standard/ Expanded	"standard"	Used Community String or Expanded Expression
		Enable	"false"	Enable of this Extended Community List
		Mode	"route"	Mode
		Route Target/ Expanded Expression	""	Match against a BGP extended community of type Route Target
		Name	""	

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
				Name of this provisioning object. Must be unique in its type
		Permit/Deny	"permit"	Action to take on a match of the Community Strings, either Permit or Deny
Gateway Port Profile	Provisioning - Gateway Profiles	Group	""	Group
		Enable	"false"	Enable flag of this provisioning object
		Enable	"false"	Enable flag for this row
		Source IP/Mask	""	Source address on the port if untagged or on the VLAN if tagged used for the outgoing BGP session
		Gateway	""	BGP Gateway referenced for this port profile
External Gateway Profile	Tenants - Gateways	Name	""	Name of this provisioning object. Must be unique in its type
	Network view - Site Gateway	Group	""	Group
		Advertisement Interval	"30"	The minimum time in seconds between sending route updates to BGP neighbor
		Anycast IP/Mask	""	The Anycast Address will be used to enable an IP routing redundancy mechanism designed to allow for transparent failover across a leaf pair at the first-hop IP router.
		Connect Timer	"120"	Time in seconds between sucessive attempts to Establish BGP session
		Default Originate	"false"	When set, switch sends a default route to BGP neighbor
		Ebgp-Multihop	"255"	Allows external BGP neighbors to establish peering session multiple network hops away.
		Enable	"false"	Enable flag of this provisioning object
		Export Route Map	""	A route-map applied to routes exported into the current tenant from the targeted BGP router with the purpose of

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
				filtering or modifying the routes
		Fabric Interconnect	"false"	
		Gateway Mode	"BGP"	Gateway Mode. Can be BGP, Static, or Default
		Helper Hop IP Address	""	Helper Hop IP Address
		Hold Timer	"180"	Time, in seconds, used to determine failure of session Keepalive messages received from remote BGP peer
		Import Route Map	""	A route-map applied to routes imported into the current tenant from the targeted BGP router with the purpose of filtering or modifying the routes
		Keepalive Timer	"60"	Interval in seconds between Keepalive messages sent to remote BGP peer
		Local AS Number	""	Local AS Number
		MD5 password	""	MD5 password
		Name	""	Name of this provisioning object. Must be unique in its type
		Neighbor AS-Number	""	Autonomous System Number of remote BGP peer
		Neighbor IP address	""	IP address of remote BGP peer
		Enable	"false"	Enable of this static route
		IPv4 Route Prefix	""	IPv4 unicast IP address followed by a subnet mask length
		Next Hop IP Address	""	Next Hop IP Address. Must be a unicast IP address
		Tenant	""	Tenant
		Source IP address	""	Source IP address used to override the default source address calculation for BGP TCP session
		Egress VLAN	""	VLAN used to carry BGP TCP session
LAG	Network View - External LAGs	Color	"anakiwa"	

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
				Choose the color to display the connectors on the network view
	Network View - on each switch	Enable	"false"	Enable flag of this provisioning object
		Eth Port Profile	""	Choose an Eth Port Profile
		Is Fabric	"false"	Indicates this LAG is used for peer-to-peer Peer-LAG/IDS link
		Is Peer Link	"false"	Indicates this LAG is used for peer-to-peer Peer-LAG/IDS link
		LACP	"true"	LACP
		Name	""	Name of this provisioning object. Must be unique in its type
Prefix List	Global Section - Route Map Definitions	Notes	""	User notes describing this provisioning object
		Enable	"false"	Enable flag of this provisioning object
		Enable	"false"	Enable of this IP Prefix List
		Greater than Equal Value	""	Match IP routes with a subnet mask greater than or equal to the value indicated
		IPv4 Prefix	""	IP address and subnet to match against
		Less than Equal Value	""	Match IP routes with a subnet mask less than or equal to the value indicated
		Permit/Deny	"permit"	Action to take on a match of the Community Strings, either Permit or Deny
		Name	""	Name of this provisioning object. Must be unique in its type
Reserved Ranges	Admin Tab - Provisioning Reserved Ranges	Base Anycast MAC Address	"00:00:00:00:01:01"	Starting MAC address to use with static Anycast gateway IP addresses
		Base BGP AS Number	"61000"	Base BGP Autonomous System Number used for switches in the fabric
		Base L3 VLAN ID	"2000"	Starting L3 VLAN ID
		Enable	"false"	Enable flag of this provisioning object

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
		Max Sites	"500"	Maximum number of Sites
		Max Switches	"2000"	Max number Switches to support in this site
		Max Tenants	"500"	Maximum number of VxLAN/ EVPN Tenants in the site
		Name	""	Name of this provisioning object. Must be unique in its type
		Paired IP Subnet	"192.168.254.0/24"	IP address range reserved for communication between paired switches
		Paired Leaf VLAN	"1999"	VLAN used to establish connection between paired leaf switches
		Paired Link VLAN	"4089"	VLAN used to establish connection between paired switches
		Router ID Base Prefix	"172.16.0.0"	Router ID starting IP address
		VTEP ID Base Prefix	"172.16.10.0"	VTEP ID starting IP address
Route Maps	Global Section - Route Map Definitions	Notes	""	User notes describing this provisioning object
		Match fields shown	""	Match fields shown
		Enable	"false"	Enable flag of this provisioning object
		Match AS Path Access List	""	Match AS Path Access List
		Match Community List	""	Match Community List
		Match EVPN Route Type	""	Match based on the indicated EVPN Route Type
		Match EVPN Route Type Default	""	Match based on the type of EVPN Route Type being Default'
		Match Extended Community List	""	Match Extended Community List
		Match Interface Number	""	Match Interface Number
		Match Interface VLAN	""	Match Interface VLAN
			""	

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
		Match IPv4 Address IP Prefix List		Match IPv4 Address IP Prefix List
		Match IPv4 Next Hop IP Prefix List	""	Match IPv4 Next Hop IP Prefix List
		Match Local Preference	""	Match BGP Local Preference value on the route
		Match Metric	""	Match Metric of the IP route entry
		Match Origin	""	Match routes based on the value of the BGP Origin attribute
		Match Peer Interface	""	Match BGP Peer port the route was learned from
		Match Peer IP Address	""	Match BGP Peer IP Address the route was learned from
		Match Peer VLAN	""	Match BGP Peer VLAN over which the route was learned
		Match Source Protocol	""	Match Routing Protocol the route originated from
		Match Tag	""	Match routes that have this value for a Tag attribute
		Match VNI	""	Match based on the VNI value
		Match VRF	""	Match VRF the route is associated with
		Name	""	Name of this provisioning object. Must be unique in its type
		Permit/Deny	"permit"	Action to take on a match of the Community Strings, either Permit or Deny
		Notes	""	User notes describing this provisioning object
		Enable	"false"	Enable flag of this provisioning object
		Name	""	Name of this provisioning object. Must be unique in its type
		Enable	"false"	Enable
		Route Map Clause	""	Route Map Clause is a collection match and set rules
Service	Tenants-Services	Anycast IP/Mask	""	Static anycast gateway address for this service

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
		DHCP Server IP	""	IP address(s) of the DHCP server for this service. May have up to four separated by commas.
		Enable	"false"	Enable flag of this provisioning object
		MTU	"1500"	MTU (Maximum Transmission Unit) The size used by a switch to determine when large packets must be broken up into smaller packets for delivery. If mismatched within a single vlan network, can cause dropped packets.
		Name	""	Name of this provisioning object. Must be unique in its type
		Tenant	""	Tenant
		VLAN	""	A Value between 1 and 4096
		VNI	"(auto)"	Indication of the outgoing VLAN layer 2 service
Static Fabric Connections	Network View - on each switch	isPredefined	"false"	Denotes this is a predefined provisioning object
		Endpoint1 for a Static Connection	""	1st Switchpoint for a Static Connection
		Endpoint2 for a Static Connection	""	2nd Switchpoint for a Static Connection
		Endpoint1 for a Static Connection	""	1st Switchpoint for a Static Connection
		Port2 for a Static Connection	""	2nd Port for a Static Connection
		Enable	"true"	Enable flag of this provisioning object
		Name	""	Name of this provisioning object. Must be unique in its type
Tenant	Tenant	DHCP Relay: Source IPs Subnet	""	Range of IP addresses (represented in IP subnet format) used to configure the source IP of each DHCP Relay on each switch that this Tenant is provisioned on.
		Enable	"true"	Enable flag of this provisioning object
		Export Route Map	""	

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
				A route-map applied to routes exported into the current tenant from other tenants with the purpose of filtering or modifying the routes
		Import Route Map	""	A route-map applied to routes imported into the current tenant from other tenants with the purpose of filtering or modifying the routes
		Name	""	Name of this provisioning object. Must be unique in its type
		Route Distinguisher	""	Route Distinguishers are used to maintain uniqueness among identical routes from different routers. If set, then routes from this Tenant will be identified with this Route Distinguisher (BGP Community). It should be two numbers separated by a colon.
		Route Target Export	""	A route-target (BGP Community) to attach while exporting routes from the current tenant. It should be a comma-separated list of BGP Communities: each Community being two numbers separated by a colon.
		Route Target Import	""	A route-target (BGP Community) to attach while importing routes into the current tenant. It should be a comma-separated list of BGP Communities: each Community being two numbers separated by a colon.
		Enable	"false"	Enable
		Tenant	""	Tenant
		Layer-3 VLAN	"(auto)"	VLAN value used to transport traffic between services of a Tenant
		Layer-3 VNI	"(auto)"	VNI value used to transport traffic between services of a Tenant
		VRF Name	"(auto)"	Virtual Routing and Forwarding instance name associated to tenants

Copyright BE Networks 2026

Parameter Group	UI View	Name	Default	Description
User Accounts	Admin Tab - Users and Permissions	isPredefined	"false"	Denotes this is a predefined provisioning object
		Base Provisioning	"true"	Capability to edit objects in the provisioning section and non-fabric lags
		Device Endpoint	"true"	Capability to edit unlocked ONT endpoints
		Device Management	"true"	Capability to manage Devices
		Enable	"false"	Enable flag of this provisioning object
		Globals	"true"	Capability to edit Badges and Radius Servers
		Idle Time Logout (min)	240"	Idle Time Logout in minutes
		Import	"true"	Capability to import
		Locked Switch Endpoint	"true"	Capability to edit locked switch endpoints
		Mod	""	Role Modifier
		Name	""	Name of this provisioning object. Must be unique in its type
		Network	"true"	Capability to edit Sites, Domains, Switch ordering, LANs, Switch Pairing, and System/Site wide management buttons
		User Permitted Sites	""	User Permitted Sites
		Role	"rw"	Role
		SD-Admin Access	"true"	Capability to access SD-Admin
		Services	"true"	Capability to edit services
		Sets	"true"	Capability to edit endpoint sets
		Site Permissions	"false"	Capability to edit Sites
		Switch Endpoint	"true"	Capability to edit unlocked switch endpoints
		Views	"true"	Capability to edit endpoint views

Copyright BE Networks 2026

6.7 END USER LICENSE AGREEMENT

READ BEFORE INSTALLING OR OPERATING THIS SOFTWARE

THIS END USER LICENSE AGREEMENT (“AGREEMENT”) IS A LEGAL AGREEMENT BETWEEN YOU, AS AN INDIVIDUAL OR ENTITY, (“LICENSEE”) AND BEYONDEDGE NETWORKS, INC., DBA BE NETWORKS, A DELAWARE CORPORATION, (“LICENSOR”) REGARDING YOUR USE OF THE BE NETWORKS FABRIC MANAGER FOR EDGE OR FABRIC MANAGER FOR CLOUD SOFTWARE (THE “SOFTWARE”). BY CLICKING “I AGREE” OR USING THE SOFTWARE, LICENSEE IS ACCEPTING AND AGREEING TO BE BOUND BY THE AGREEMENT (THIS “AGREEMENT”). THE SOFTWARE CONSISTS OF CERTAIN LICENSOR-DEVELOPED SOFTWARE AS WELL AS CERTAIN SOFTWARE DEVELOPED BY THIRD PARTIES. THE SOFTWARE IS LICENSED (NOT SOLD) TO YOU. THIS AGREEMENT SUPERSEDES ANY PREVIOUS VERSIONS. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, DO NOT USE THE SOFTWARE.

6.7.1 1. DEFINITIONS

a) “Capacity” means specific size criteria, e.g., managed port or hardware device count. Software License details regarding Capacity and Licensee current usage information is available through the Orchestration Platform user interface. b) “Partner” means an authorized Licensor partner that in their ordinary course of business resells the Software to its Customers and is not a Managed Service Provider. The Partner is responsible for providing Customer with pre- and post-sales activities including network design, configuration, provisioning and Tier 1 and Tier 2 support. The Partner is also responsible for transferring the title of any hardware and the License purchased from Licensor. c) “Partner Customers” means Licensees who purchased the Software from the Partner. d) “Documentation” means all instructions and other documentation related to the Software provided by Licensor. Licensee may make a copy of the Documentation for archival purposes. e) "End User" means any commercial customer who purchases Software for its internal use and not for resale. The End User, or the entity that the End User represents, is the Licensee. f) “License File” means the file provided by Licensor specifically to the target Orchestration Platform and is viewable by Licensee on the Orchestration Platform user interface that specifies the licensed parameters of the Software, e.g., licensed capacity, current usage, and contacts for licensing and support. g) “Managed Service Provider (MSP)” means a Licensee that in their ordinary course of business provides value-added services to its customers by managing the operation of the Software. h) “Managed Service Provider (MSP) Customers” mean customers of Managed Service Providers. i) “Normal Business Hours” means the period of time commencing at 0800 hours and running until 1700 hours on the same day, according to local time in Dallas, Texas. j) “Orchestration Platform” means the Software component (Virtual Network Commander – vNetC) which provides web browser access to the Software for the Licensee’s management and configuration of its installed instances of the Software and for tracking License status. k) “Subscription Term” means the period of time, so long as the Licensee complies with all the terms of the Agreement, beginning on the Install Date to the end date at which Licensee ceases to utilize the Software licensed hereunder. l) “Tier 1 Support” is the initial support level responsible for basic Licensee issues Tier 1 Support gathers relevant information and determines the Licensee’s issue by analyzing the symptoms and diagnosing the root cause of the underlying problem. m) “Tier 2 Support” is responsible for assisting Tier I Support personnel in solving more difficult technical problems by confirming the validity of the problem and trying known solutions to these more complex issues. Tier 2 Support escalation to Tier 3 Support presumes either a software bug or a limitation of the troubleshooting tools. n) “Tier 3 Support” denotes expert level troubleshooting and analysis methods and presumes the Software has a bug or a limitation in the troubleshooting tools. Tier 3 Support first validates the Tier 2 Support escalation to Tier 3 Support.

6.7.2 2. LICENSE STRUCTURE, ACTIVATION AND SUBSCRIPTION FEE SCHEDULE

a) License Grant. Provided that Licensee complies with all terms and conditions of this Agreement, including the payment to Licensor or its authorized Partner, as identified on the License File, all Subscription Fees, Licensor grants Licensee a subscription-based, non-exclusive, non-transferable (except as expressly set forth herein), limited license during the applicable Subscription Term (as defined below) to download and use the Software, solely for the Licensee’s own internal purposes or for managing the functionality of the Software for its customers, (“MSP Customers”) in the ordinary course of Licensee’s business. b) Documentation. During the Subscription Term, Licensor grants Licensee the right to use and copy the Documentation in connection with Licensee’s authorized use of the Software provided hereunder; provided all copies of the Documentation contain the same notices of copyright and ownership as the original. c) License Structure. Licenses are attached to the Media Access Control (MAC) address of the server on which the Orchestration Platform Software instance is installed. Software is licensed on a Capacity as specified in purchase order, master purchase agreement or other valid purchasing instrument. d) License Activation. To enable the Software for use by Licensee, Licensee must download and install the Software on an acceptable virtual machine(s), as specified by Licensor, in the Licensee’s on- premises or cloud-based computing environment. Upon receipt of a valid purchase order for the Software, Licensor will provide Licensee with the necessary license keys to enable full operational functionality of the Software. e) Subscription Activation and Billing. The Software is licensed on a subscription basis (“Subscription”). The Billing Period shall be as specified in purchase order, master purchase agreement or other valid purchasing instrument (the “Billing Period”). For purposes of this Agreement, the delivery date for the Software will be the date on which the License File is downloaded to the

Copyright BE Networks 2026

Orchestration Platform (“the Install Date”). The Install Date defines the first day of the Billing Period. Subscription Fees are due in advance of the beginning of the next Billing Period. f) Fees, the Software automatically prevents the Licensee from making any changes (e.g., services provisioning or hardware device additions or changeouts) to the underlying network. Reinstatement of a valid License requires the payment in full of all in-arrears Subscription Fees and a one-time reinstatement fee billed as twenty-five (25) percent of the aggregate in-arrears Subscription Fee. g) License Adds During Billing Period. If additional Licenses are purchased during a Billing Period, the end of that Billing Period shall be adjusted such that all Subscription Fees (including the new ones) Nonpayment of Subscription Fees. In the event Licensee fails to stay current with Subscription will be due on the same date. The adjusted date is calculated by (1) calculating the total amount of time that the newly purchased Licenses extend past the current Billing Period, and by (2) redistributing that dollar-time interval across all Licenses, both new and existing, on a pro-rata basis. The Licensee may also extend the License validity date by providing a purchase order for the dollar amount accrued for a specified time interval for which the Billing Period is to be extended. Evaluation License. Any Software licensed for evaluation purposes or other limited, temporary h) use as authorized by Licensor (“Evaluation Software”), Licensee use of the Evaluation Software is only permitted for the period limited by the license key or otherwise stated by Licensor in writing. The Evaluation Software is licensed “AS-IS” without support or warranty of any kind, expressed or implied. Licensor does not assume any liability arising from any use of the Evaluation Software. Licensee may not publish any results of benchmark tests run on the Evaluation Software without first obtaining written approval from Licensor. Licensee authorizes Licensor to use any feedback Licensee provides Licensor in connection with Licensee’s use of the Evaluation Software.

6.7.3 3. AGREEMENT TERM.

This Agreement expires on the day that Licensee ceases to use the Software licensed hereunder (the “Term”).

6.7.4 4. COPYRIGHT; RESTRICTIONS.

As between the Licensee and Licensor, all right, title and interest in and to the Software shall remain in Licensor. The Software provided hereunder is licensed, not sold, under this Agreement by Licensor for use only under the terms set forth herein. The Software is the proprietary and confidential property of Licensor, and the Software (including all intellectual property rights therein) is exclusively owned by Licensor and as may be applicable, Licensor’s licensors. This Agreement does not grant Licensee any rights in or to the source code of the Software. All rights not expressly granted to Licensee in this Agreement are reserved to Licensor. Licensee shall not, directly or indirectly, reverse engineer or aid in the reverse engineering of all or any part of the Software. “Reverse engineering” includes, without limitation, decompiling, disassembly, sniffing, or otherwise deriving source code. Licensee may not transfer any System provided hereunder to a third party without the prior written approval of Licensor. ANY UNAUTHORIZED USE OF THE SOFTWARE PROVIDED HEREUNDER WILL AUTOMATICALLY VOID THE LICENSE GRANT SET FORTH IN SECTION 2(A). LICENSEE UNDERSTANDS AND AGREES THAT VIOLATION OF THIS SECTION WOULD CAUSE IRREPARABLE DAMAGE TO LICENSOR’S BUSINESS AND MAY SUBJECT LICENSEE TO CLAIMS BY LICENSOR FOR COPYRIGHT INFRINGEMENT, INJUNCTIVE RELIEF, AND/OR MONETARY DAMAGES.

6.7.5 5. THIRD-PARTY SOFTWARE

Certain software that is provided to Licensor by third parties may be delivered as an integrated component of the Software and is subject to the licenses of their respective owners (“Third Party Software”). This Agreement shall apply to all Third-Party Software providers as if they were Licensor and to the Third-Party Software as if it were the Software.

6.7.6 6. MAINTENANCE AND SUPPORT SERVICES.

Provided Licensee is current with all Subscription Fees, Licensee is entitled to receive without additional charge, all maintenance and upgrade releases of the Software as Licensor, in its sole discretion, makes them generally available. Licensor does not provide any support services for the Software under this EULA except Tier 3 Support to Partners or MSPs. Any additional BE Networks support services to Partners, MSPs, or Licensees shall be as described in a separate agreement or other legal document between the parties.

6.7.7 7. LIMITED WARRANTY.

Licensor warrants that during the Term and contingent on Licensee being current with all Subscription Fees, for Licensee’s benefit alone (and not for the benefit of any Licensee Customer), the Software delivered under this Agreement will perform substantially in accordance with the specifications published by Licensor. If, during the Term, Licensee notifies Licensor in writing of a material defect in the Software and if Licensor is able to replicate the defect, then Licensor will work to correct the defect or provide Licensee with replacement Software that corrects the material defect. Any replacement will be warranted for the remainder of the Term. EXCEPT AS EXPRESSLY SET FORTH ABOVE, THE SOFTWARE IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, AND NO OTHER WARRANTIES, EITHER EXPRESSED OR IMPLIED, ARE MADE WITH RESPECT TO A SYSTEM, INCLUDING (BUT NOT LIMITED TO) ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE OR NON-INFRINGEMENT, OR ANY OTHER WARRANTIES THAT MAY ARISE FROM USAGE OF TRADE OR COURSE OF DEALING. LICENSOR DOES NOT

Copyright BE Networks 2026

WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE OF OR THE RESULTS OF THE USE OF A SYSTEM IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE AND DOES NOT WARRANT THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED OR ERROR FREE OR THAT LICENSOR WILL BE ABLE TO CORRECT ALL DEFECTS. LICENSOR EXPRESSLY DISCLAIMS ANY WARRANTIES NOT STATED HEREIN.

6.7.8 8. REMEDIES.

Licensor's entire liability and Licensee's exclusive remedy under the above limited warranty shall be that Licensor, in its sole discretion, will use reasonable commercial efforts to correct or provide a workaround or fix for the Software for any reproducible defect in the Software (each such action a "Remedy"), at Licensor's cost and expense, provided that Licensee gives written notice itemizing such material defect to Licensor during the Term. Because of the nature of computer software, however, Licensor cannot guarantee that any Remedy nor any other services or support provided under this Agreement will permanently cure any Software defect. The limited warranty is void if failure of the Software has resulted from Licensee or any Licensee Customer's improper use.

6.7.9 9. LIMITATION OF LIABILITY.

In no event shall either party have the right to recover damages of any nature in connection with the business contemplated by this Agreement due to the expiration or permitted or lawful termination of this Agreement. Each party waives and releases the other from any claim to compensation or indemnity for termination of the license rights hereunder unless such termination is in breach of this AGREEMENT. IN NO EVENT SHALL LICENSOR BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ANY CLAIM THAT ARISES FROM OR RELATES IN ANY WAY TO THIS AGREEMENT OR USE OF THE SOFTWARE PROVIDED HEREUNDER, REGARDLESS OF THE FORM OF ACTION ALLEGED, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, AND NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED OR EXCLUSIVE REMEDY OF ANY KIND. IF LICENSOR SHALL BE LIABLE TO LICENSEE, THEN IN NO EVENT AND REGARDLESS OF THE FORM OF ACTION ALLEGED, SHALL LICENSEE BE ENTITLED TO ANY MONETARY DAMAGES AGAINST LICENSOR IN EXCESS OF THE SOFTWARE FEES LICENSEE PAID TO LICENSOR FOR LICENSEE CUSTOMER(S) GIVING RISE TO THE CLAIM. Licensee understands that the applicable fees reflect the allocation of risk set forth in this Agreement and that Licensor would not have made any Software available for Licensee's use without the limitations of liability set forth in this Agreement.

6.7.10 10. GOVERNMENT USE; PROCUREMENT.

If the Software is being licensed under the terms of a proposal or agreement with the U.S. Government or any contractor on its behalf, the Software is commercial computer software, is developed exclusively at private expense, and (a) if acquired by or on behalf of a civilian agency, shall be subject to the terms of this commercial computer software license as specified in 48 C.F.R. 12.212 of the Federal Acquisition Regulations and its successors; and (b) if acquired by or on behalf of units of the Department of Defense ("DOD") shall be subject to the terms of this commercial computer software license as specified in 48 C.F.R. 227.7202-2, DOD FAR Supplement and its successors. If the Software will be provided or made available to the U.S. Government, any use, duplication, or disclosure by the U.S. Government of the Software shall be subject to the restrictions applicable to proprietary commercial computer software set forth in subparagraph (c)(1)(ii) of The Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 or subparagraphs (c)(1) and (2) of the Commercial Computer Software - Restricted Rights clause at 48 C.F.R. 52.227-19, as applicable. The contractor/manufacturer is BE Networks, 2600 N. Central Expressway, Suite 950, Richardson, TX 75080.

6.7.11 11. TERMINATION

This Agreement may be terminated as follows: (a) immediately without notice upon Licensee's insolvency, assignment to creditors, appointment of a trustee or commencement of a proceeding seeking bankruptcy; (b) by Licensor in the event of a material breach of this Agreement by Licensee that is not cured to the reasonable satisfaction of Licensor within thirty (30) days of providing written notice of such material breach; and (c) upon written notice by either party at least sixty (60) days prior to the end of the then current Term. Upon termination of this Agreement for any reason, Licensee will not be entitled to any refund of any fees paid hereunder. Upon any termination of this Agreement, Licensee must destroy any and all copies of the Documentation and immediately cease all use of the Software together with any copies thereof and provide written certification by an officer of your company that the foregoing requirements have been complied with; and (iii) all terms and conditions of this Agreement shall cease, except Sections 4, 7, 9, 12, and 18.

6.7.12 12. CONFIDENTIALITY; TRADE SECRETS.

The Software and Documentation contain proprietary and confidential information of, and trade secrets owned by, Licensor and its licensors. Such trade secrets include, without limitation, the specific design, structure and logic of individual Software programs, their interactions with other portions of the Software, both internal and external, the programming techniques employed therein, and instructions on its use. You agree to hold the

Copyright BE Networks 2026

Software, Documentation and all other confidential or proprietary information provided to you by Licensor (“Confidential Information”) in strict confidence and not to use or disclose such Confidential Information in any way except as expressly as specified in Section 14. Licensee agrees to protect the Confidential Information at least to the same extent that you protect your similar confidential information, but in no event with less than reasonable care. There can be no adequate remedy at law for any breach of your obligations hereunder and upon any such breach, or threat thereof, Licensor shall be entitled, without the necessity of posting any bond, to injunctions and other appropriate equitable relief in addition to whatever remedies it may have at law.

6.7.13 13. INDEMNIFICATION.

Licensee shall defend, indemnify, and hold Licensor and its respective officers, directors, employees, representatives and agents harmless from and against any and all claims, demands, damages, losses, liabilities, costs or expenses (including, without limitation, reasonable attorney fees) (“Claims”) arising from or incurred in connection with any of the following: (a) Licensee’s breach of this Agreement; or (b) any and all Claims by or on behalf of Licensee’s End Users. Licensor agrees to indemnify, defend and hold harmless Licensee from and against any and all Claims resulting from or arising out of allegations that any Software or the use thereof infringe upon, misappropriate or violate any United States patents, copyrights, or trade secret rights of persons, firms or entities who are not parties to this Agreement; provided that Licensee (a) promptly notifies Licensor, in writing, of any Claim of such alleged infringement or misappropriation involving the Systems of which it becomes aware, and (b) permits Licensor to control the defense settlement, adjustment or compromise of any such Claim using counsel of its choice. Unless Licensor fails to perform its obligations pursuant to this section, Licensee shall have no authority to settle any Claim on behalf of Licensor.

6.7.14 14. ASSIGNMENT.

Licensee may assign this agreement to an affiliate of Licensee or to any successor of Licensee by reason of merger, reorganization, sale of all or substantially all of the assets, change of control or operation of law. Except as set forth above, Licensee may not assign or in any manner transfer any of its rights or obligations under this Agreement to any third party without the prior written approval of Licensor, which consent shall not be unreasonably withheld or delayed. Licensor may assign any of its rights or obligations under this Agreement to a third party without the prior approval of Licensee.

6.7.15 15. EXPORT RESTRICTIONS.

Licensee acknowledges that the Software provided hereunder are of U.S. origin and may be subject to export regulations by the U.S. government. Licensee agrees to comply with all applicable international and national laws that apply to the Software, including the U.S. Export Administration Regulations, as well as end-user, end-use and destination restrictions issued by the U.S. and other governments.

6.7.16 16. CHANGES AND ENTIRETY.

Except as specifically provided herein, no provision of this Agreement may be waived or modified unless agreed to in writing by both parties. Neither party shall be deemed to have waived any right to or remedy under this Agreement unless such waiver is expressed in writing. This Agreement (including the attachments listed on the Cover Page) constitutes the entire agreement between the parties and supersedes all prior written or oral agreements between the parties relating to the subject matter stated herein.

6.7.17 17. SEVERABILITY AND GOVERNANCE.

If any provision in this Agreement shall for any reason be held invalid, void, or unenforceable, the remaining provisions shall remain in full force. Any provisions of this Agreement prohibited by the law of any state shall be ineffective to the extent of such prohibition without effecting the remaining provisions of this Agreement. Licensor is a Delaware corporation, and this Agreement is governed by the laws of the State of Texas, without regard to any rule or principle that would refer to or apply the substantive law of another state or jurisdiction. Each party by its respective signature acknowledges that it has read and understands this Agreement; that this Agreement constitutes the entire agreement, understandings, and representations, expressed or implied, between Licensee and Licensor with respect to the Software license and that this Agreement supersedes all prior communications between the parties, including all oral or written proposals.

BE Networks End User License Agreement (EULA), Rev. 3 February 2026

Copyright BE Networks 2026

Changelog **Version Updates**

Copyright BE Networks 2026

7. Changelog

7.1 Verity 6.7.0

7.1.1 General Updates

Feature	Description
IVN-24712 REST API Authentication	The ML service REST API (`/api/v1`) now requires API key authentication. API keys and the admin secret are generated at install and kept across upgrades. The Grafana Infinity datasource is provisioned with the API key, so dashboards authenticate automatically
IVN-24712 Grafana Hardening	Anonymous access to datasource configuration is blocked (401 on read and write). Session cookies are now Secure and SameSite. Login rate-limiting is enabled. The Grafana admin password is generated at install instead of using a default
IVN-24712 Mutual TLS Between vNetC and Monitoring	Both vNetC and the monitoring VM verify each other's certificates. Before this release, only the server side was verified
IVN-24712 Automatic TLS Provisioning	Install creates a CA for each deployment (in `/be_install/satori_ca/`) and issues a server certificate, with no manual certificate steps. The gateway serves this certificate over HTTPS on port 8443. The CA can only sign certificates for the monitoring VM's own IP and hostname (`nameConstraints` , `pathlen:0`), so a stolen CA key can't be used for any other host. Server certificates are reissued automatically when they are within 30 days of expiry or the IP changes, and vNetC keeps trusting them. The CA is only regenerated if the hostname changes or regeneration is forced, which requires re-pinning the CA on vNetC
IVN-24712 vNetC Certificate Trust (Outbound)	The monitoring VM checks vNetC's certificate against a pinned CA for announce, NBI, and package download calls. `trust_vnetc_ca.sh` fetches the certificate chain from `VNETC_ROOT_URL` , validates it, and installs only the CA. `NBI_VERIFY_SSL` now points to this CA (`/app/certs/vnetc.pem`) for every outbound HTTPS call. The CA fingerprint is confirmed once, either at an interactive prompt or by setting `VNETC_CA_FINGERPRINT` before an unattended install. An unattended install with no pinned fingerprint stops with an error. The pin is stored in `/be_install/vnetc_ca.fingerprint` and is kept across upgrades. Trust is re-checked on every upgrade: a new vNetC server certificate is accepted automatically, but a changed CA is refused unless forced
IVN-24712 Monitoring Certificate Trust on vNetC (Inbound)	vNetC must trust the monitoring CA before it can connect to the gateway on port 8443. Copy `satori_ca.pem` , `satori_ca.fingerprint` , and `trust_satori_ca_on_vnetc.sh` to the vNetC host and run the script. Alternatively, run the script with `--fetch` to download the CA and confirm its fingerprint. The script installs the CA in `/var/ns/ca-certs-loaded/` and pins its fingerprint in `/var/ns/etc/satori_ca.fingerprint` . The CA is included in vNetC backups and copied to HA peers automatically. If the monitoring CA is regenerated, run the script again with force to re-pin it
IVN-24712 Data Store Authentication	Redis now requires a password generated at install and no longer exposes a host port. Neo4j credentials are generated at install. APOC is limited to four required procedures. The Neo4j browser route is removed from the gateway
IVN-24712 Container Hardening	Monitoring services now run with reduced privileges, following container security best practices. Access between services and the host is restricted, and gateway routing is tightened.
IVN-24712 Pinned Container Images	Container images are pinned to exact versions instead of floating tags. Builds are reproducible and can't silently pick up unvetted upstream changes
IVN-24712 Install Manifest Verification	Install checks the image and payload manifest before deploying. A tampered bundle fails early
IVN-24712 Credential Cleanup	Hardcoded credentials are removed from the environment files that ship with the product. Those values are now empty placeholders
IVN-24712 Documentation Updates	The monitoring installation and configuration guides now cover automatic TLS setup, key generation, and the updated install flow

Copyright BE Networks 2026

7.2 Verity 6.6.1

7.2.1 General Updates

Feature	Description
Branding Enhancements	New branding options and other Product Change Requirements
MCP Server - Phase 1	Not integrated with AI/OPS
LDAP for SONiC	Manage LDAP in SONiC Switches
Dell OS10 SNMP Collector & Dashboards	Gather telemetry from existing OS10 devices.

7.3 Verity 6.6

7.3.1 New Features

Feature	Description
PLM-375 Update to wolfSSL Commercial Release v5.9.1 for 432a-v2 using the FIPSv2 (Cert 3389)	Addresses CVE-2026-5194, a high-severity ECDSA verification vulnerability affecting X.509, TLS, and DTLS use cases; upgrades to wolfSSL 5.9.1 FIPS v2 (Cert 3389) on the 432a-v2 platform
PLM-369 Prometheus/Grafana External Alerting	Allows the user to extend alerting and metrics to external monitoring systems
PLM-361 Multiple Fabric Support	Support multiple datacenter instances in one orchestration deployment. Fabric Tree Navigation and Switch management improvements
PLM-335 Monitoring Dashboards 6.6	Executive Summary, Port Health, Compact Switch Dashboards, and general improvements
PLM-334 GUI: Tree Browser for Fabric/Devices	Introduces a tree-style browser interface for managing fabrics and devices within the GUI
PLM-333 Extensible Telemetry (OME Support)	Add metrics for physical servers connected to Dell OME servers
PLM-329 Tenancy tree browser	Introduces a tree browser UI for navigating tenancy hierarchy
PLM-316 Spectrum-X Reference Architecture 2.0 for NVIDIA/Cumulus	Support for Reference Architecture 2.0
PLM-306 Manage NTP, SYSLOG, DNS, TACACS (and more) in SONiC Switches	Manage miscellaneous operation features of SONiC switches
PLM-298 Hostname/ARP/ND Cache Workbench	Enhance connected user lookup to include Hostnames and IP address along with MAC address searches
PLM-294 AI-Ops Chatbot - Phase 2: Change Sets, Multi-Turn and Multilingual Conversations	Adds ability for natural language chat based provisioning changes as well as supporting multi-lingual exchanges
PLM-276 DHCP with MAC Binding	Assign IP addresses based on expected discovered devices assigned Static IP addresses

Copyright BE Networks 2026

7.3.2 Resolved Issues and Minor Feature Additions

Issue	Description
IVN-21354	Add BGP Option to remove all private AS numbers in advertised AS route paths
IVN-22039	Add BGP Gateways configured with "allowas-in" option
IVN-22047	Add Thresholding Alerts for Verity Components
IVN-22112	Add MTU support to the Eth Port Setting
IVN-22392	Unify device controller form and switchpoint form
IVN-22481	Add Metric for Link Transitions
IVN-23095	Add PFC and ECN Counters per port to switch and port dashboards
IVN-23234	IPV6 Unnumbered Tenant and Fabric Gateways
IVN-23261	Add Port Setting Option for "standalone-link-training"
IVN-23262	Add option to adjust IP attached-host advertise to layer 3 Service

Copyright BE Networks 2026

8. L2 Edge

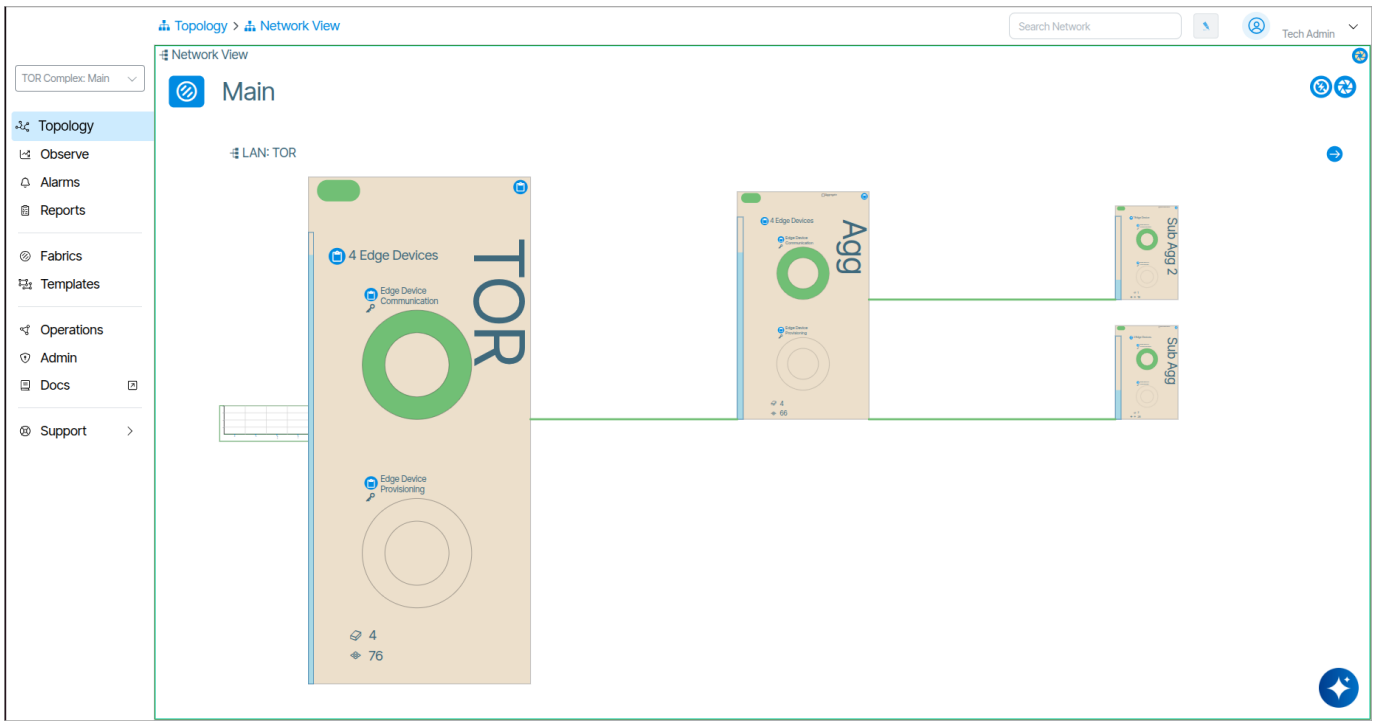
8.1 Verity L2 Edge

Introduction

Verity is available in one of two different versions: **L2 Edge** and **Data Center**. The **L2 Edge** solution is historically referred to as **Enterprise**.

As the two versions evolve, the **L2 Edge** documentation covers only its deviations from the **Data Center** documentation, deferring to it where applicable.

The Verity L2 Edge system is designed to interface with a Campus Architecture. The network topology is auto-discovered by the Verity software. Additionally, no manual configuration is required for interconnects between White Box devices — TORs, AGGs, and ONTs. This allows for simple replacement of devices and movement of port connections in the event of hardware component or connection failures. Black Box devices that support LLDP are also automatically discoverable.



8.1.1 L2 Edge Installation

For installation details see Installation Details

Copyright BE Networks 2026

8.2 Manual Onboarding

8.2.1 Before You Begin

- The security lockdown process requires all equipment used in the process are located on the same network subnet as the file server containing any files needed for the lockdown process. For switch and ONT devices, that server must reside on the PC/host that is physically connected to the device.
- The user must provide SSL Certificates that are installed during the lock down process. Certificates are contained in a file format called PEM (defined in RFC 1422), which can contain either just the public certificate (e.g., CA certificate files in /etc./SSL/certs) or a full certificate chain with the private key, public key, and root certificates.
- **Certificate files may have different names depending upon the entity generating them**. Also, in the case of many switches and ONT devices, the files usually have serialized names so they can be tracked. For the purposes of this document the following names and descriptions are used:

1. **crt.pem** – file used by devices and the SDLC to validate the orchestration platform certificate
2. **client.pem** – file used by devices and the SDLC to identify themselves
3. **client-cert.pem** – file used by the SDLC/ACS to validate clients. Sometimes referred to as the certificate-chain file
4. **server.pem** - used by the vNETC and ACS to identify themselves to the managed devices/SDLC. This is the “web server key and certificate” file
5. **DoD root certificate chain files** - multiple files loaded via drag/drop to "Root Certificate Chains"
6. **External Authentication Client Certificate** - loaded via drag/drop (use for client connection to the RADIUS server
7. **Certificate Revocation List** - single CRL file loaded via drag/drop

- The following SW loads are required to support Security Lockdown
- Core – 6.3.0275
- Firmware Package 1.6.1.57

> **Device Requirements**
>
> The devices (switches and ONTs) MUST BE IN A CLEAN FACTORY RESET CONDITION before starting the process or connecting to any system. If they are connected to a system before starting the process, once they contain a valid load, the user should factory reset the device appropriately.

> **Certificate Names**
>
> Certificate names are used throughout this document; however, as mentioned, the certificate names you work with may not be the same as the examples provided.

SSH

To configure any device for lockdown, SSH must be used to access the device command line interface. The setup process assumes the user has requisite knowledge of SSH.

What File Transfer Protocol and Software Do I Need for Certificates?

Installing certificates can be done with HTTP or TFTP and examples are provided for both HTTP and TFTP. It is the user’s responsibility to choose and manage their file transfer server application.

Command Line Writing Convention

Users are prompted to input commands throughout the procedure. The commands are shown as bold text. The following example prompts the user to use the command line interface and apply the commands **admin** and **security**.

Copyright BE Networks 2026

Main/ admin/security

The result is the command line directory is changed to **Security Base Menu**. Text in a grey background represents the content displayed on the command line window.

```
Customer Menu

ACS Security Menu Commands

keys -\> SSH keys get, del & show Menu
cert -\> SSL certificate get, del & show Menu
host -\> Host ip/name add, del & show Menu
? -- Help or ? followed by the command string for specific help.
Q \<- One Level Back
E . Exit

Main/ administration/ Security Base Menu/
```

Device Command Line QuickStart

This section of the documentation provides an overview of the device command line software interface, which consists of menus and commands.

In the example below, the command to the left of each -\ > symbol navigates to the menu that is to the right of each -> symbol.

```
Customer Menu

VCF -\> Vendor Configuration Menu
CONN -\> Connection Information
CRAFT -\> Craft Menu
DNLD -\> Software/Download Menu
```

Menu symbols may have different forms such as:

-\ >

\ >\ >

--

.

In the following example, the command **CRAFT** is associated with the **Craft Menu**. To navigate to the **Craft Menu**, type **CRAFT** and press Enter.

```
Customer Menu

VCF -\> Vendor Configuration Menu
CONN -\> Connection Information
CRAFT -\> Craft Menu
DNLD -\> Software/Download Menu
FAULT -\> Fault Menu
INV -- Display Inventory
OMCI -\> OMCI Menu
TRACE -\> Trace Menu
VP -\> Voice Processing Menu
NET -\> Network
RG -\> Residential Gateway
DBG -\> Debug menu

? -- Help or ? followed by the command string for specific help.

E . EXIT

Enter Option(s) Value: CRAFT
```

Basic Navigation

To navigate to a previous menu, you use the **Q** command.

Some windows prompt you with \< > icons. These icons convey that user defined inputs are to be submitted. User defined inputs are called parameters. The following code prompts you to type a URL of your choosing.

```
<url\>
```


Copyright BE Networks 2026

In the following example the **get-cert** command requires between 1 to 3 parameters. The parameters include a certificate type, URL, and a file name.

The certificate type is prefaced by the **-t** or **–type** command. 

```
Enter Option(s) Value: get-cert

Command requires from 1 to 3 parameters

Usage:

get-cert \<-t\|--type\> \<cert type\> \<url\>

Usage:

\-t, --type \<crt\|client\> - Certificate type (required)

\<url\> - In the form "http://10.12.99.1/\<...\>/\<filename\>
```

```
get-cert \<-t\|--type\> \<cert type\> \<url\>
```

8.2.2 Security lockdown for ONT Device

The security lockdown process for an ONT device is summarized in four steps:

1. Log into the device using SSH
2. Install certificates
3. Lock down the device
4. Repeat this for all ONT devices that are connecting to the locked down system.

Connecting to the Device

HARDWARE SETUP

Connect an ethernet cable from a PC to ethernet port #1 on the ONT.

Using SSH to Access the Device Interface

NETWORK CONFIGURATION

The ONT has a backdoor IP address of 10.12.99.32. The PC must be in the same subnet, such as 10.12.99.30. The IP address of the PC's Ethernet Network Interface Card (NIC) must be set to 10.12.99.30.

AUTHENTICATION

SSH must be used to access and configure the device. It is assumed that the user has the requisite knowledge to use SSH.

1. **SSH**: 10.12.99.32
2. **Login**: tech
3. **Password**: \<enter default tech password>
4. Upon successful authentication the interface window will be displayed, and the **Craft Menu** is available. 

About SSL Certificates

ONT REQUIRES TWO SSL CERTIFICATES

The ONT requires *two* certificates - **crt.pem** and **client.pem** - to be present. The **crt.pem** is used to validate the server certificate of the ACS. The **client.pem** is the certificate sent by the ONT to authenticate the ONT when it is talking to the ACS and vNETC.

Certificates Have Unique Names

The certificates you install on your system will have unique names. For this document, they are referred to as **crt.pem** and **client.pem.**

Copyright BE Networks 2026

8.2.3 Installing the ONT CRT Certificate

To begin, navigate to the SSL certificate Base Menu

Customer Menu/CRAFT/SEC

```
Customer Menu -\> Craft Menu -\> Security Base Menu -\> SSL certificate Base Menu
```

The **SSL certificate Base Menu** provides an option to add, delete, list, and display SSL certificates content.

1. Select **cert**.
2. Select **get-cert -t \<cert type> \<url>**.
3. The window that appears prompts you for the certificate type, the URL and file name.
4. Submit the command. The result of your entry is formatted like the following examples depending on the transfer protocol you use:

HTTP

```
get-cert -t crt <http://10.12.99.30/crt.pem>
```

TFTP

```
get-cert -t crt tftp://10.12.99.30/crt.pem
```

VERIFY SSL CERTIFICATE DOWNLOAD

1. Use the **show-cert -t crt** command to verify that the SSL Certificate process completed properly.
2. The certificate contents are presented.

8.2.4 Installing the ONT Client Certificate

1. Navigate to the SSL certificate Base Menu **Customer Menu/CRAFT/SEC**
2. The **SSL certificate Base Menu** provides an option to add, delete, list, and display SSL certificates content.
3. Select **cert**.
4. Select **get-cert -t \<cert type> \<url>**.
5. The window prompts you for the certificate type, the URL and file name.

```
Enter Option(s) Value: get-cert

Command requires from 1 to 3 parameters

Usage:

get-cert \<-t\|--type\> \<cert type\> \<url\>

Usage:

\-t, --type \<crt\|client\> - Certificate type (required)

\<url\> - In the form "http://10.12.99.32/\<...\>/\<filename\>
```

Submit the command. The result of your entry is formatted like the following examples depending on the transfer protocol you use.

HTTP

```
get-cert -t client <http://10.12.99.30/client.pem>
```

TFTP

```
get-cert -t client tftp://10.12.99.30/client.pem
```

VERIFY SSL CERTIFICATE DOWNLOAD

1. Use the **show-cert -t client** command to verify that the SSL Certificate process completed properly.
2. The certificate contents are presented.

Copyright BE Networks 2026

8.2.5 Lock-Down Mode

Setting the device into **Lock-Down Mode** is the last step in the configuration.

LOCKING DOWN

1. Navigate to the **Security Base Menu**.
2. Select **lock** (Type **q** and press **Enter** to go back one level).

REBOOT

This is the last step in preparation to place the ONT in lockdown mode. 1. Navigate to the **Craft Menu** (type **q** and press **Enter** to go back one level). 1. Type **r** and press **Enter**.

VERIFY LOCKDOWN

1. Use SSH to authenticate into the device.
2. If you are unable to authenticate the lockdown is successful.

8.3 Security Lockdown for Switches

The lockdown process for a Switch device can be summarized in four steps:

1. Log into the device using SSH
2. Install certificates
3. Lock down device
4. Repeat this section for all switching devices that are connecting to the locked down system.

> **Before You Begin**
>
> If the switch is going to be installed as a “Top of Rack” (TOR) and the uplink connection does not support using LLDP for automatic TOR and Uplink detection, please refer to the switch documentation to manually set up the TOR before proceeding.

8.3.1 Securing the Device

Use SSH to Access the Device Interface

NETWORK CONFIGURATION

The switch has a backdoor IP address of 10.12.99.32. The PC must be in the same subnet, such as 10.12.99.30. The IP address of the PC's NIC must be set to 10.12.99.30.

AUTHENTICATION

SSH must be used to access and configure the device. It is assumed that the user has the requisite knowledge to use SSH.

1. **SSH**: 10.12.99.32
2. **Login**: tech
3. **Password**: \<enter default tech password>
4. Upon successful authentication the interface window is displayed with the prompt **Main#**. 

About SSL Certificates

SWITCH REQUIRES TWO SSL CERTIFICATES

The switch requires *two* certificates - **crt.pem** and **client.pem** - to be present. The **crt.pem** is used to validate the server certificate of the vNETC. The **client.pem** is the certificate sent by the Switch to authenticate the Switch when it is talking to the vNETC.

Copyright BE Networks 2026

What File Transfer Protocol Do I Need for Certificates?

Installing certificates can be done with either HTTP or TFTP.

8.3.2 Installing the CRT Certificate

1. Navigate to the SSL certificate Base Menu **admin/security/cert**.
2. The **SSL certificate Base Menu** provides an option to add, delete, list, and display SSL certificates content.
3. Select **cert**.
4. **get-cert -t crt \<url/filename>**
5. The window that appears prompts you for the certificate type, the URL, and the file name.
6. Submit the command.
7. The result of your entry is formatted like the following examples depending on the transfer protocol you use:

HTTP

```
get-cert -t crt <http://10.12.99.30/crt.pem>
```

TFTP

```
get-cert -t crt tftp://10.12.99.30/crt.pem
```

VERIFY SSL CERTIFICATE DOWNLOAD

1. Use the **show-cert --type crt** command to verify that the SSL Certificate process completed properly.
2. The certificate contents are presented.

8.3.3 Installing the Client Certificate

1. Navigate to the SSL certificate Base Menu **Main/admin/security/cert**.
2. The **SSL certificate Base Menu** provides an option to add, delete, list, and display SSL certificates content.
3. Select **cert**.
4. Select **get-cert -t client \<tftp://10.12.99.32/2S846000025.pem>**.
5. The window that appears prompts you for the certificate type, the URL, and the file name.

```
Enter Option(s) Value: get-cert

Command requires from 1 to 3 parameters

Usage:

get-cert \<-t\|--type\> \<cert type\> \<url\>

Usage:

\-t, --type \<crt\|client\> - Certificate type (required)

\<url\> - In the form "http://10.12.99.32/\<...\>/\<filename\>
```

6. Submit the command. The result of your entry is formatted like the following examples depending on the transfer protocol you use.

HTTP

```
get-cert -t client <http://10.12.99.30/client.pem>
```

TFTP

```
get-cert -t client tftp://10.12.99.30/client.pem
```

VERIFY SSL CERTIFICATE DOWNLOAD

1. Use the **show-cert --type client** command to verify that the SSL Certificate process completed properly.

Copyright BE Networks 2026

2. The certificate contents are presented.

8.3.4 Lock-Down Mode

Enter q to go back a level in the menu to set the device into **Lock-Down** mode.

LOCKING DOWN

1. Navigate to **Main/admin/security**
2. Select **lock**.

REBOOT

This is the last step in preparation to place the Switch in lockdown mode. Enter q to go back a level

1. Navigate to the **Main Menu**.
2. Type **reboot** and press **Enter**.

VERIFY LOCKDOWN

Use SSH to attempt to access the device. If you are unable to login using the same credentials the lockdown is successful.

8.4 Security Lockdown for SDLC/ACS Virtual Machine Applications

The configuration process can be summarized in four steps.

1. Log into the SDLC application using SSH. The user's PC and file server containing certificates (may be the same machine) must have access to the management VLAN that the SDLC is connected to.
2. Install certificates for the SDLC and ACS application in the SDLC
3. Lock down the SDLC application
4. After the vNETC is locked down and the SDLC is subsequently connected, the user will create the ACS application and it inherits the certificates and lockdown status from the SDLC. There is no explicit process for the ACS during this step of the procedure.

8.4.1 Securing the SDLC Application

Use SSH to Access the Device Interface

SSH must be used to access and configure the device. It is assumed that the user has the requisite knowledge to use SSH.

Login: tech

Password: \<enter tech default password>

Upon successful authentication the interface window is displayed with the prompt **Main#**.

Install SSL Certificates

SDLC REQUIRES TWO SSL CERTIFICATES - CRT.PEM AND CLIENT.PEM - TO BE PRESENT.

8.4.2 Installing the CRT Certificate

To begin, navigate to the SSL certificate Base Menu.

/ Main / admin / security / cert

```
Main/ administration/ Security Base Menu/ SSL certificate Base Menu\#
```

The **SSL certificate Base Menu** provides an option to add, delete, list, and display SSL certificates content.

Copyright BE Networks 2026

Select **cert.**

Select **get-cert -I br4094 -t crt \<url/filename>**.

```
get-cert -- Get certificate file from specified URL
del-cert -- Delete certificate from persistent memory
list-certs -- List certificates from persistent memory
show-cert -- Show certificate from persistent memory
? -- Help or ? followed by the command string for specific help.
Q \<- One Level Back
E . Exit
```

The window that appears prompts you for the certificate type, the URL, and the file name.

```
Enter Option(s) Value: get-cert

Command requires from 1 to 3 parameters

Usage:

get-cert \<-t\|--type\> \<cert type\> \<url\>

Usage:

\-t, --type \<crt\|client\> - Certificate type (required)

\<url\> - In the form "http://10.12.99.1/\<...\>/\<filename\>
```

Submit the command. The result of your entry is formatted like the following examples depending on the transfer protocol you use.

HTTP

```
get-cert -i br4094 -t crt <http://10.12.99.30/crt.pem>
```

TFTP

```
get-cert -i br4094 -t crt tftp://10.12.99.30/crt.pem
```

VERIFY SSL CERTIFICATE DOWNLOAD

1. Use the **show-cert -t crt** command to verify that the SSL Certificate process completed properly.
2. The certificate contents are presented.

8.4.3 Installing the Client Certificate

To begin, navigate to the SSL certificate Base Menu.

/ Main / admin / security / cert

```
/ Main / administration / Security Base Menu / SSL certificate Base Menu \#
```

The **SSL certificate Base Menu** provides an option to add, delete, list, and display SSL certificates content.

Select **cert.**

Select **get-cert -i br4094 -t client \<url/filename>**.

```
get-cert -- Get certificate file from specified URL
del-cert -- Delete certificate from persistent memory
list-certs -- List certificates from persistent memory
show-cert -- Show certificate from persistent memory
? -- Help or ? followed by the command string for specific help.
Q \<- One Level Back
E . Exit
```

The window that appears prompts you for the certificate type, the URL, and the file name.

```
Enter Option(s) Value: get-cert

Command requires from 1 to 3 parameters

Usage:

get-cert \<-t\> \<cert type\> \<url\>
```


Copyright BE Networks 2026

```
Usage:

\-t, -type \<crt\|client\> - Certificate type (required)

\<url\> - In the form "http://10.12.99.1/\<...\>/\<filename\>
```

Submit the command. The result of your entry is formatted like the following examples depending on the transfer protocol you use.

HTTP

```
get-cert -i br4094 -t client <http://10.12.99.30/client.pem>
```

TFTP

```
get-cert -t client tftp://10.12.99.30/client.pem
```

VERIFY SSL CERTIFICATE DOWNLOAD

1. Use the **show-cert -type client** command to verify that the SSL Certificate process completed properly.
2. The certificate contents are presented.

8.4.4 Securing the ACS Application

Certificate Inheritance

The ACS inherits the certificates from the SDLC when the ACS is created.**

Install SSL Certificates

The ACS requires four certificates: a **server.pem** and **client-crt.pem** for ONT communication, as well as a **client.pem** and **crt.pem** for vNETC communication.

8.4.5 Installing the CRT Certificate for ONT Communication

1. To begin, navigate to the ACS SSL certificate Base Menu.
2. Enter **q** to return to Security Menu

acs/cert

```
/ Main / administration / Security Base Menu / ACS Security Menu / SSL certificate Base Menu \#
```

1. The **SSL certificate Base Menu** provides an option to add, delete, list, and display SSL certificates content.
2. Select **get-cert -i br4094 -t crt \<url/filename>**.

```
get-cert -- Get certificate file from specified URL
del-cert -- Delete certificate from persistent memory
list-certs -- List certificates from persistent memory
show-cert -- Show certificate from persistent memory
? -- Help or ? followed by the command string for specific help.
Q \<- One Level Back
E . Exit
```

The window that appears prompts you for the certificate type, the URL, and the file name.

```
Command requires from 1 to 3 parameters

Usage:

get-cert \<-t\> \<cert type\> \<url\>

Usage:

\-t, --type \<crt\|client\> - Certificate type (required)
```


Copyright BE Networks 2026

```
\<url\> - In the form "http://10.12.99.1/\<...\>/\<filename\>
```

Submit the command. The result of your entry is formatted like the following example.

```
**get-cert -i br4094 -t client-crt http://10.101.1.4/download/ca-devices.pem**
```

VERIFY SSL CERTIFICATE DOWNLOAD

1. Use the **show-cert -t client-crt** command to verify that the SSL Certificate process completed properly.
2. The certificate contents are presented.

8.4.6 Installing the Client Certificate for ONT Communication

```
Main/ administration/ Security Base Menu/ ACS Security Menu/ SSL certificate Base Menu\#
```

1. The **SSL certificate Base Menu** provides an option to add, delete, list, and display SSL certificates content.
2. Select **get-cert -t server \<url/filename>**.
3. The window that appears prompts you for the certificate type, the URL, and the file name.

```
Command requires from 1 to 3 parameters

Usage:

get-cert \<-t\> \<cert type\> \<url\>

Usage:

\-t, --type \<crt\|client\> - Certificate type (required)

\<url\> - In the form "http://10.12.99.1/\<...\>/\<filename\>
```

Submit the command. The result of your entry is formatted like the following example.

```
**get-cert -i br4094 -t server http://10.101.1.4/download/server.pem**
```

VERIFY SSL CERTIFICATE DOWNLOAD

1. Use the **show-cert -t server** command to verify that the SSL Certificate process completed properly.
2. The certificate contents are presented.

8.4.7 Lock-Down Mode

Setting the device into **Lock-Down** mode is the last step in the configuration:

1. Navigate to **Main/ administration/ Security Base Menu#** or enter **q** twice to return to Main Security Menu
2. Select **lock**.

REBOOT

This is the last step in preparation to place the SDLC in lockdown mode. After this step, the SDLC should be rebooted for the lockdown to be in effect.

VERIFY LOCKDOWN

Use SSH to authenticate into the SDLC. If you are unable to authenticate the lockdown is successful.

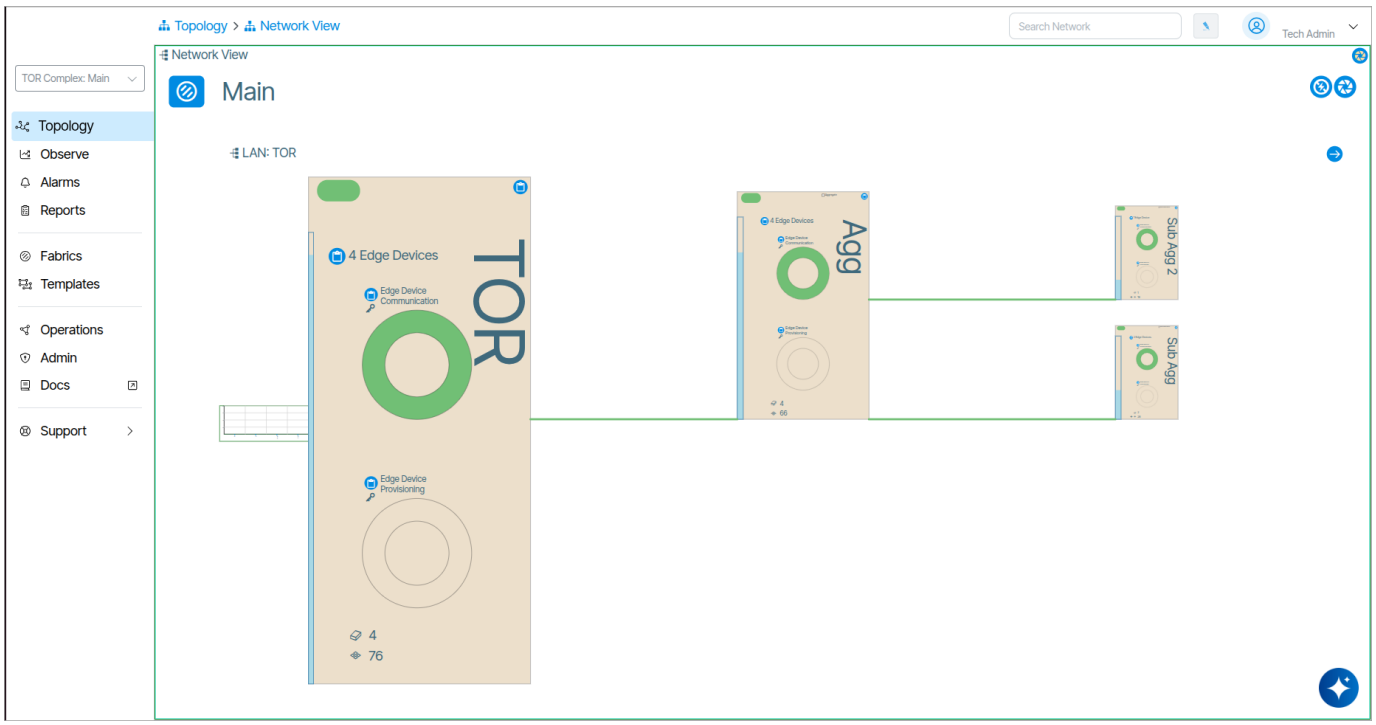
> ⚠ **SDLC Power State**
>
> POWER DOWN THE SDLC application until the vNETC lockdown is completed.

Copyright BE Networks 2026

8.5 Application Reference

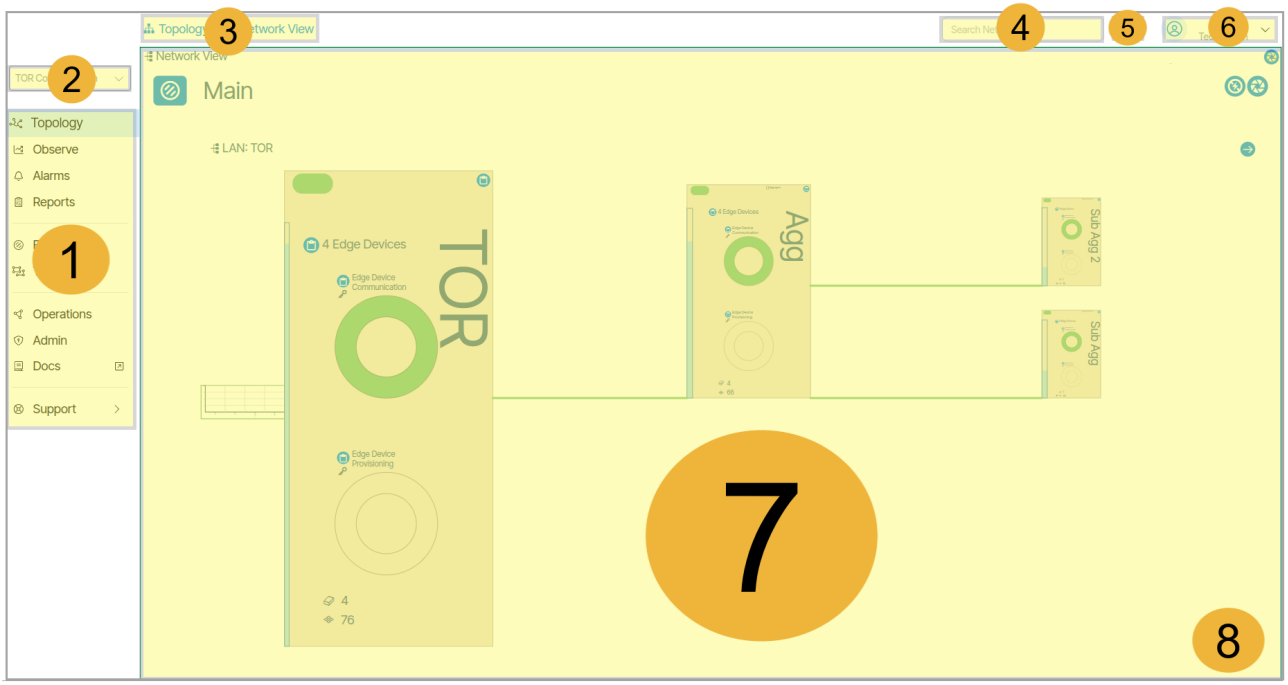
8.5.1 UI Navigation

Verity L2 Edge Workspace



The Verity L2 Edge Workspace is your single pane of glass for setting up services, profiles, and service bundle templates, and applying them to the network's service endpoints. From here, you can drill into each section to add, edit, or delete LAN elements.

Copyright BE Networks 2026



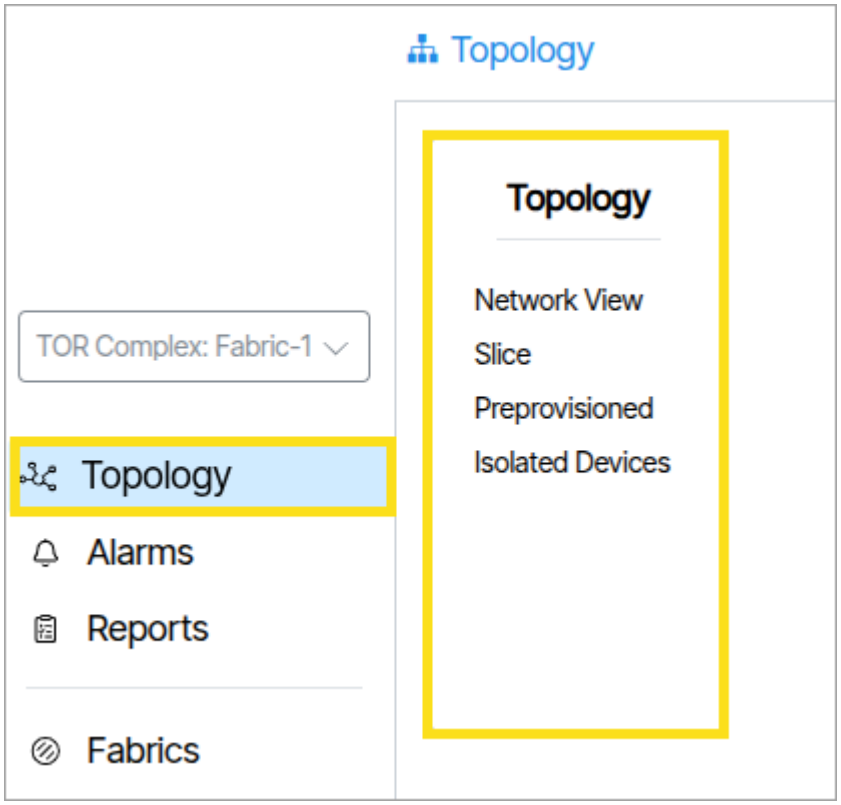
1. Main Menu
2. Site Selection
3. Navigation Depth Tracker
4. Search
5. Highlight
6. User Info & Sign Out
7. Workspace
8. AI-Ops launcher

Copyright BE Networks 2026

8.5.2 Topology

On L2 Edge systems **Topology** contains links to:

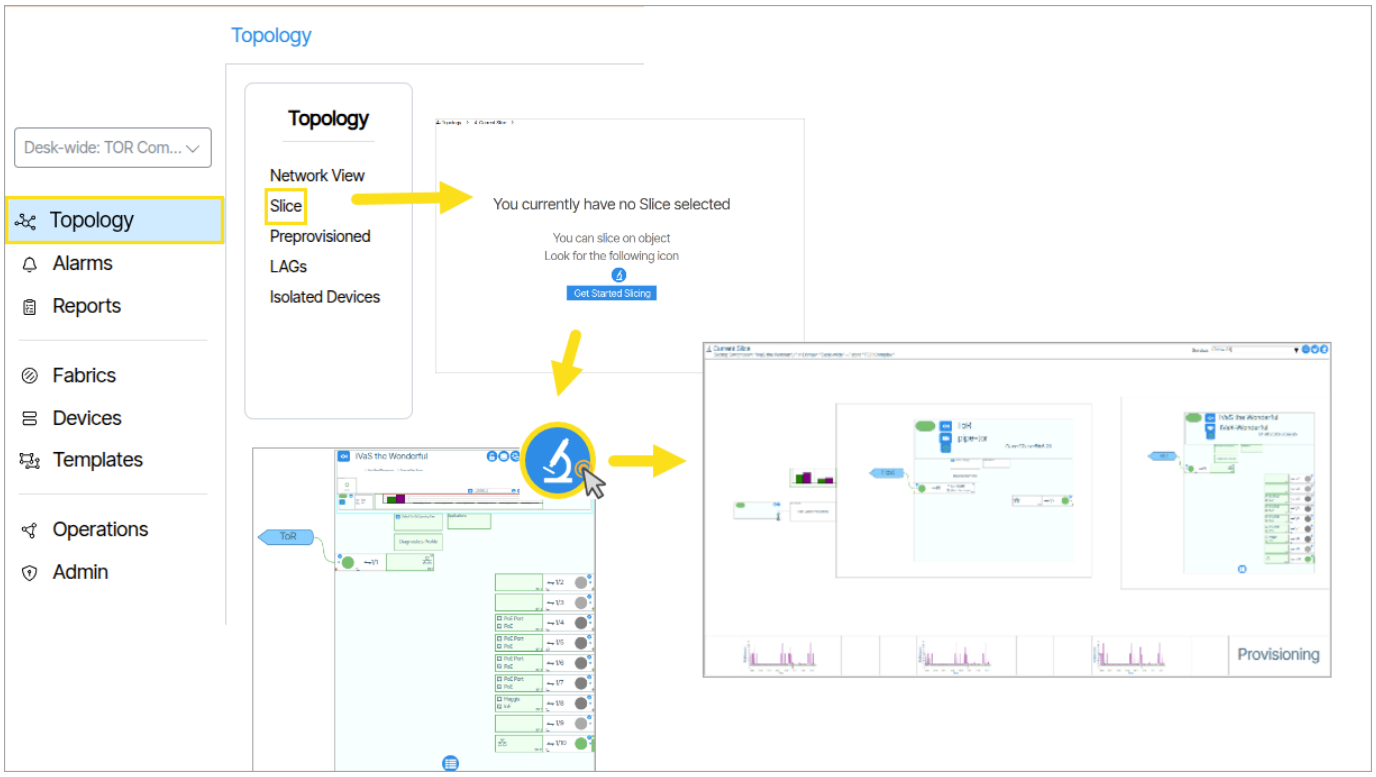
- Network View
- Slice
- Preprovisioned
- LAGs
- Isolated Devices



Slice

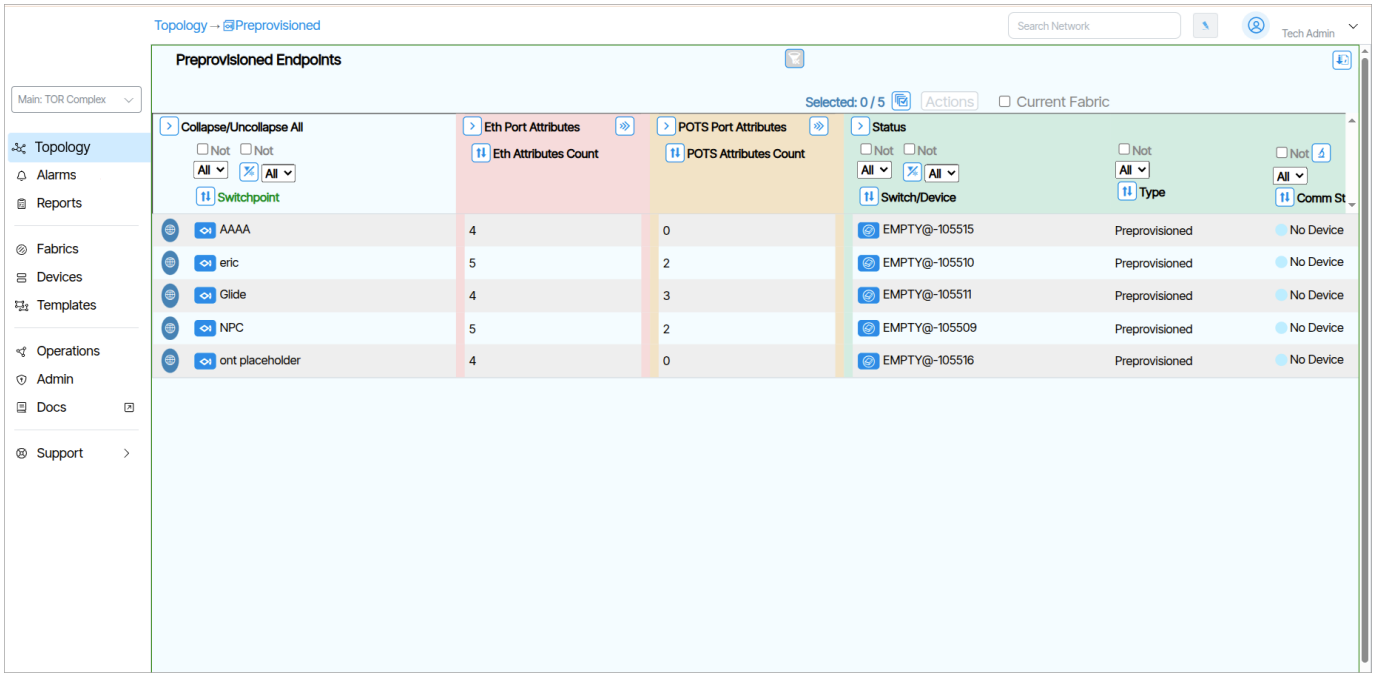
Slice view is a troubleshooting tool that displays the end-to-end view of services being delivered to the ports on the service Endpoint device. It shows the connectivity between all the network components used as well as the relevant provisioning objects to deliver the service. Zooming into the various components gives you more detail about that component. Optical levels are displayed for the optical connections between devices. Essentially, if there is a problem delivering a service to an end user port, the information leading to the disruption is in this single pane. MAC workbench and system logs can be interrogated directly from the slice view focusing on activity relevant to the devices in the slice.

Copyright BE Networks 2026



Preprovisioned

Opens a report that lists preprovisioned ONTs.



LAGS

LAG is an acronym for Link Aggregation Group. Link aggregation is a popular method for grouping multiple ethernet connections into a single logical link.

More information: LAGS

Copyright BE Networks 2026

Isolated Devices

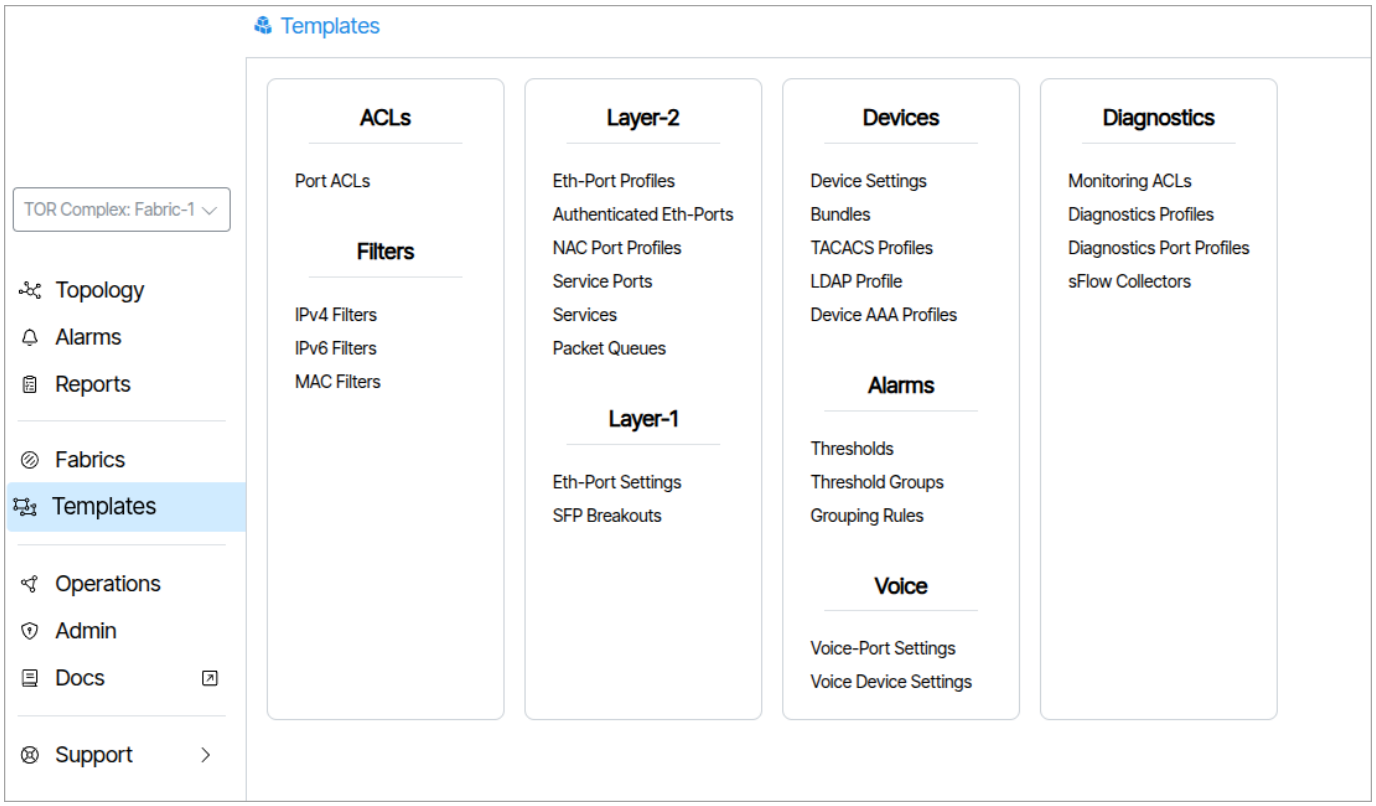
This report list devices not connected to the fabric. These are ONTs that are managed but are not connected to any TOR through present or past connections.

More information: Topology

Copyright BE Networks 2026

8.5.3 Templates

The **Templates** section is composed of tools used to configure and create reusable network device settings and policies.

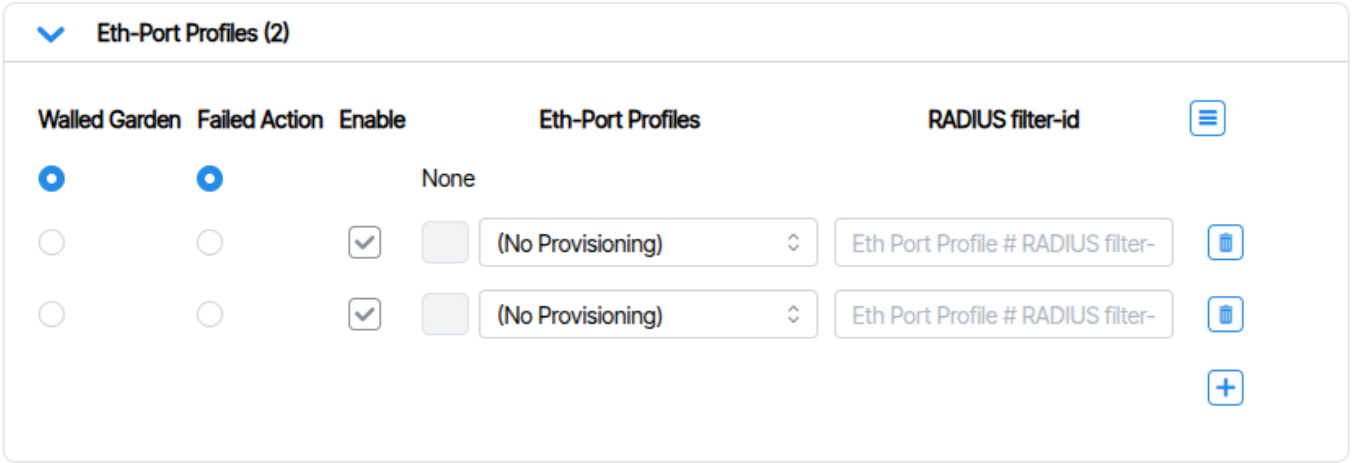


More Information: Templates

Authenticated Eth-Ports

An Authenticated Ethernet Port profile is comprised of a list of Ethernet Port Profiles. This feature is used on networks that employ user detection and authentication via 802.1x protocols. The Managed End Device communicates to connected user devices as they connect to the assigned port, and relay control messages to a Radius server, to validate the user and select and enable the appropriate Ethernet Port Profile on the port.

The “Settings” section allows a change in the connection mode the user device contains, ranging from 3 different modes. Port Mode, which is the standard mode, Single Client Mode which disables traffic from a second client, only allowing the authenticated client’s traffic to pass. Lastly, Multiple Client Mode which allows multiple clients’ traffic to pass.



Copyright BE Networks 2026

Walled Garden: A restricted network that a device gets dropped into when it hasn't authenticated (or can't). Typically it only allows access to basic features like a login portal or a remediation server, not the real network. "(No Provisioning)" here means no walled garden is set up and there no fallback network.

Failed Action: Specifies how the port responds when authentication fails. When set to "None", no additional action is taken and the device is denied network access. Other options may redirect the device to a restricted network, such as the walled garden.

Enable: Turns row on or off without deleting it.

Eth-Port Profiles: The Ethernet Port Profile this row points to. "(No Provisioning)" means nothing has been picked yet, so the row is effectively empty.

RADIUS filter-id: This is a matching key. When a user authenticates, the RADIUS server sends back a Filter-Id value (a standard RADIUS attribute, just a text string). The switch compares that string against this column, and whichever row matches is the Ethernet Port Profile that gets applied to the port. In this way, the Filter-Id maps an authenticated user's identity to the Ethernet Port Profile assigned to their port.

Authentication Workflow Summary

When a device connects to the port, the switch forwards its credentials to the RADIUS server for validation. Upon successful authentication, the RADIUS server returns a Filter-Id attribute. The switch matches this value against the RADIUS filter-id entries in the profile and applies the corresponding Eth-Port Profile to the port. If authentication fails, the Failed Action and Walled Garden settings determine how the port handles the device.

Eth-Port-Profiles (n)

An **Eth-Port-Profile** contains:

1. A collection of Services. Each Service has its own VLAN settings, plus optional ingress/egress ACLs and a MAC filter for that Service's traffic.
2. Profile-level settings that apply to the whole port: - Default Port ACL (ingress and egress) - Details: Port Monitoring and Trusted Port - Port Labeling: a label and icon - TLS

Service Ports

Service Ports are normally assigned to links or LAGs connected to the TOR switch of the network. The user can edit the maximum downstream rate allowed on a per service basis. Other parameters such as spanning tree and port monitoring are set on the physical link.

Bundles

A **Bundle** is a group of settings used to configure, program and set the state of a switch or ONT.

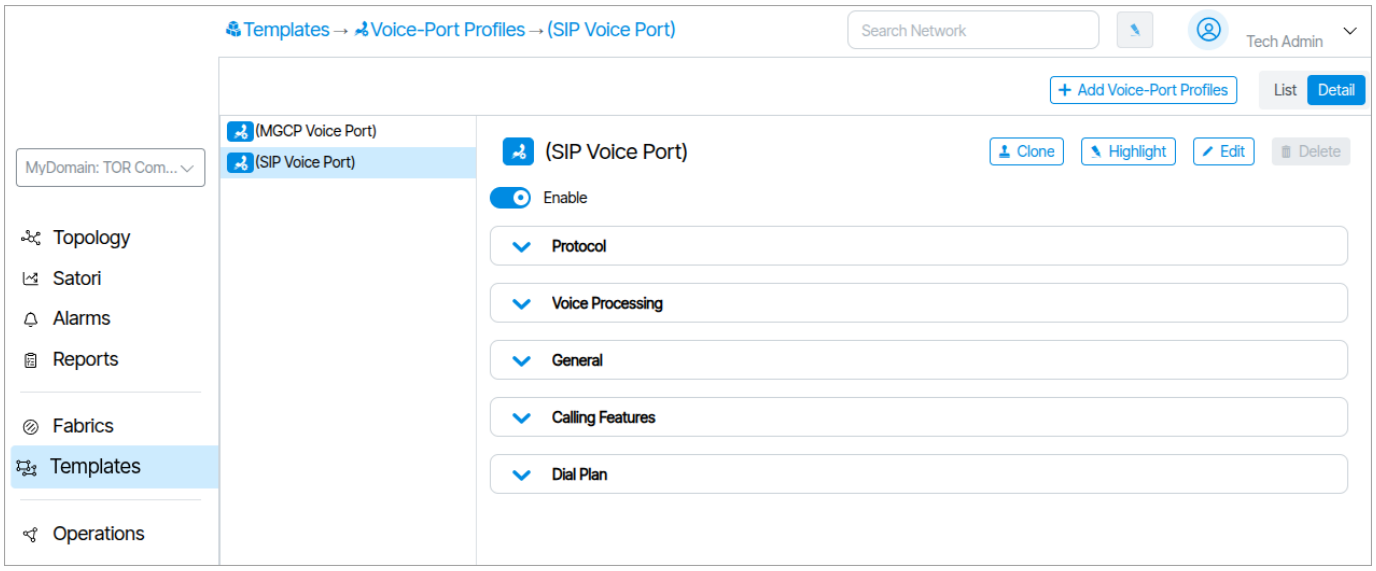
For more information read: How to Create a Bundle.

Voice Port Settings

Voice Port Settings contain parameters related to VoIP protocols, interfaces and related features. An example of a Voice Port Setting is shown below. It covers various voice settings and features that are port specific.

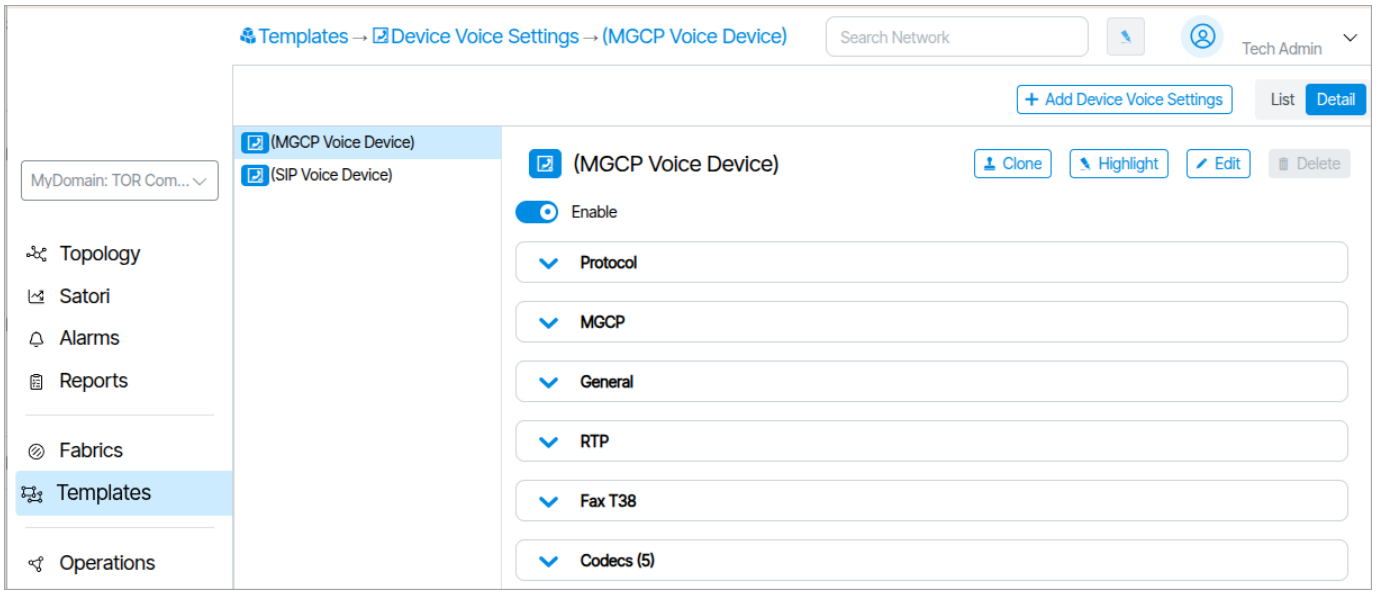
Voice Port Settings can be accessed from within the rendered ONT in Topology ().

Copyright BE Networks 2026



Voice Device Settings

Voice Device Settings are the settings used for configuring the VOIP stack on the ONT. There are two formats for the protocol field based on whether SIP or MGCP is used. Voice Device Settings can be accessed from within the rendered ONT in Topology ( ).



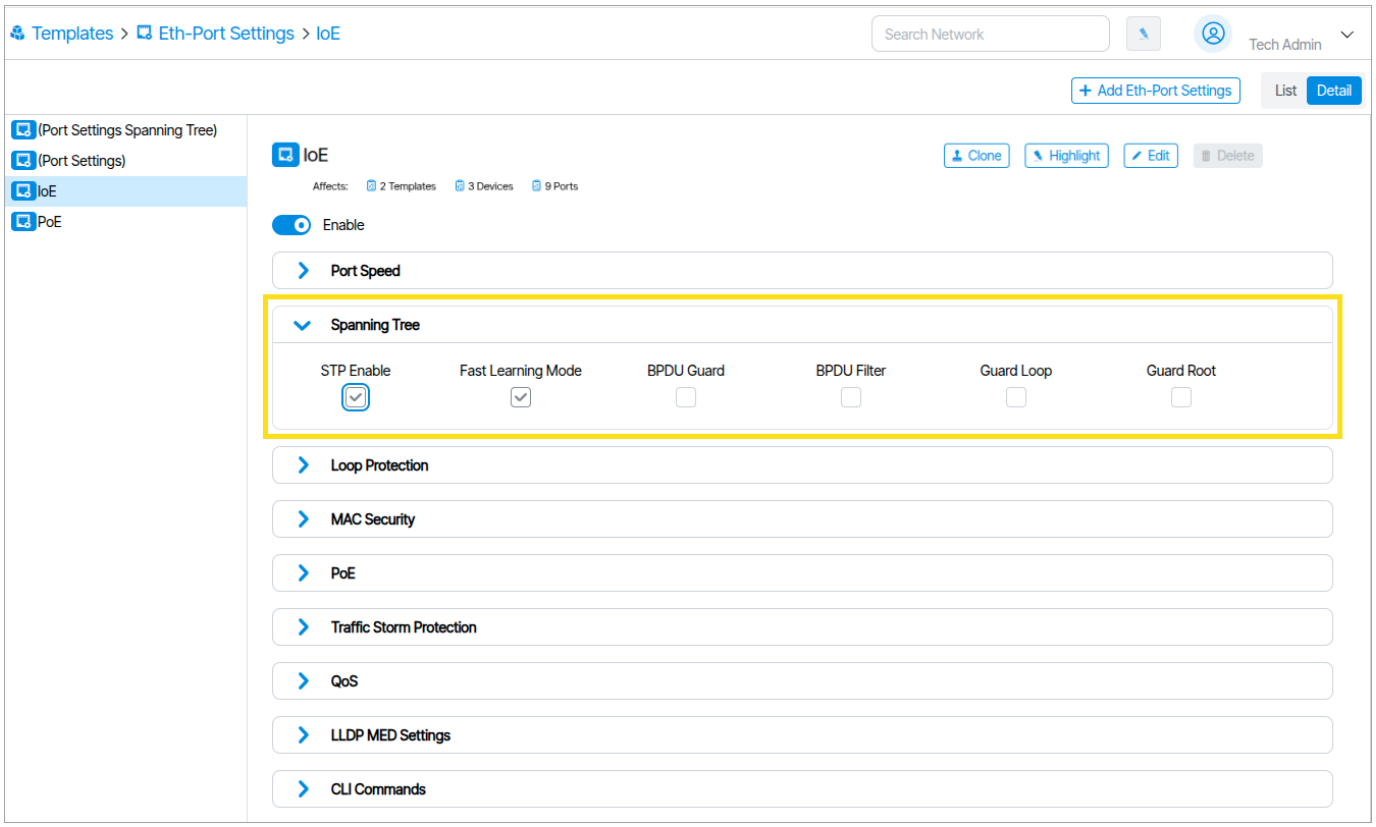
Eth-Port-Settings

An **Eth Port Setting** is a reusable profile that defines how a physical Ethernet port behaves at Layer 1 and Layer 2.

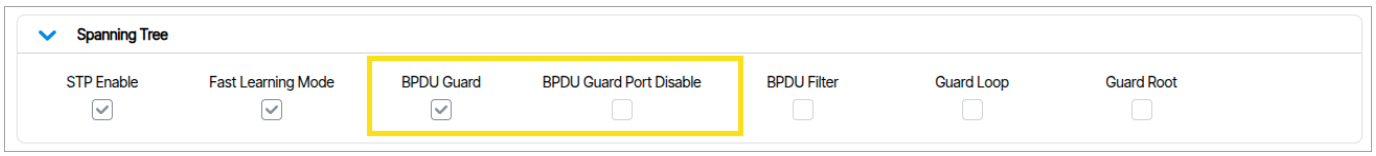
SPANNING TREE

Spanning Tree (STP) is a protocol that stops network loops when switches are connected in a circle. Spanning Tree parameters are configurable in **Eth-Port-Settings**.

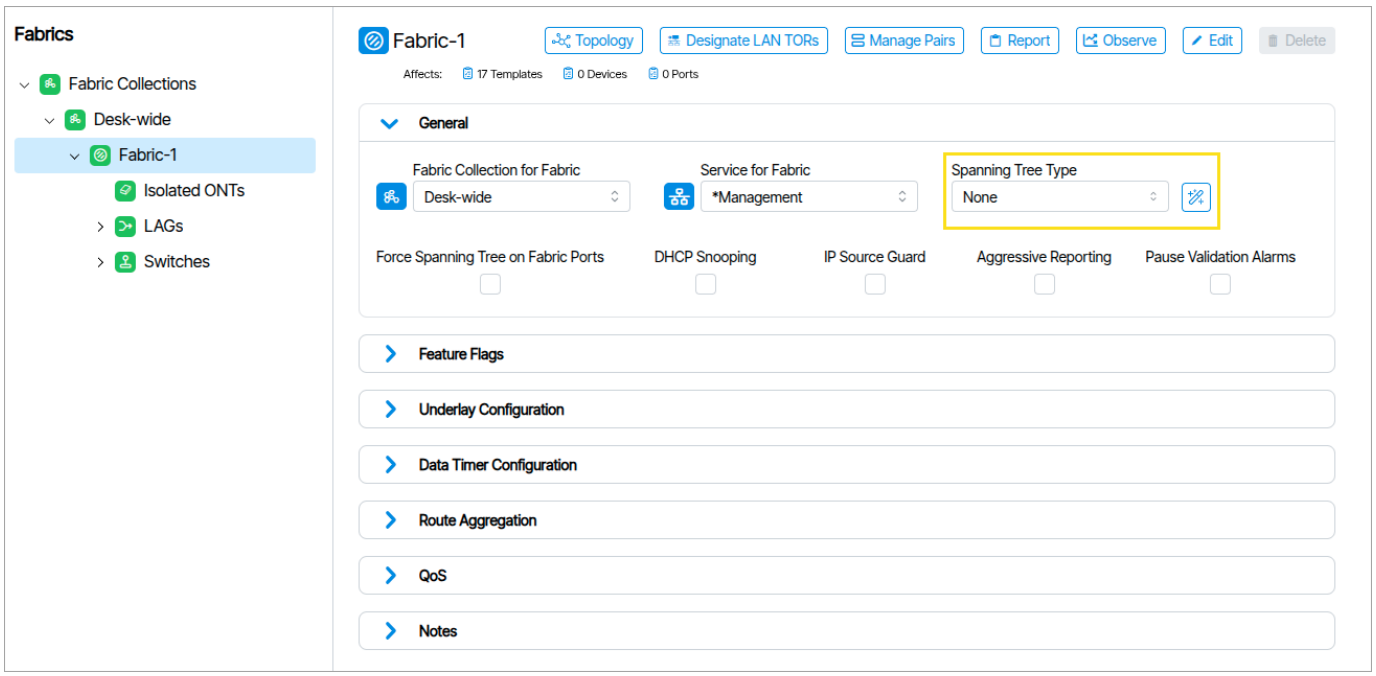
Copyright BE Networks 2026



Checking **STP Enable** and **BPDU Guard** reveals the **BPDU Guard Port Disable** option.



The Spanning Tree type (VLAN, PORT, MST) is controlled in the Fabrics Settings.

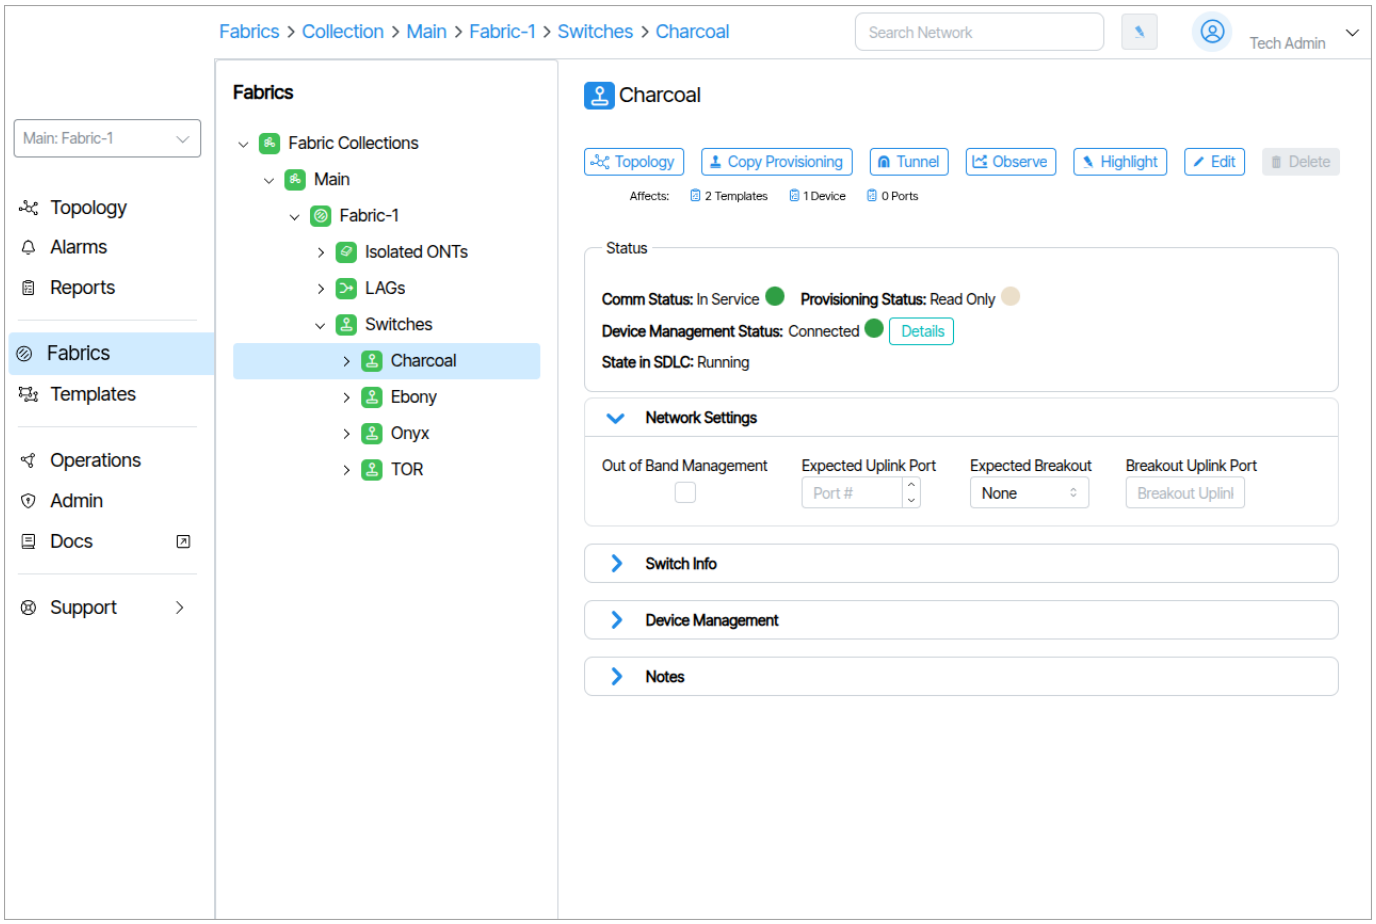


Copyright BE Networks 2026

8.5.4 Fabrics

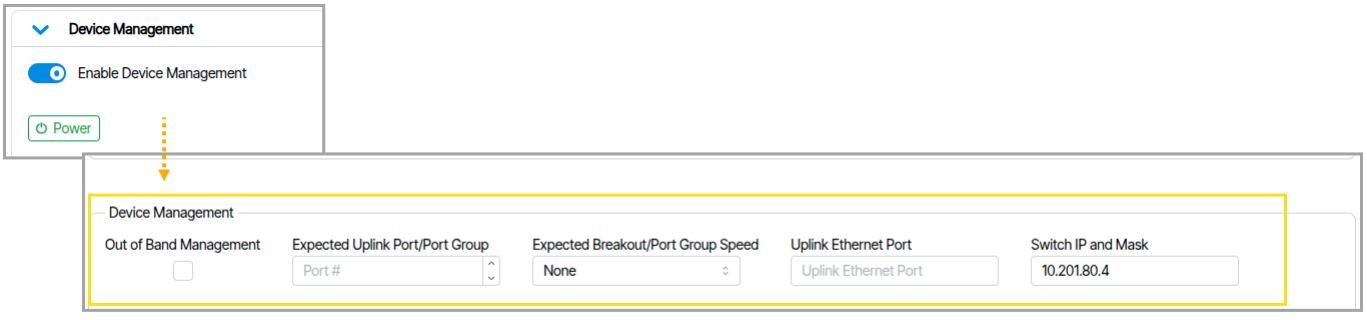
Fabrics are isolated, self-contained arrangements of network interconnections that form a topology, including the switches, links, and logical structure.

More Information: Fabrics



Device Settings (Breakout Configuration Settings)

Management Port Configuration



Problem: ZTP fails on non-copper TOR/aggregation switches because the initial `configdb.json` provisions management VLAN on the first 32 ports — but these switches typically use higher-speed uplink ports, often in breakout configuration.

Copyright BE Networks 2026

Solution: Port configuration fields are added to the switchpoint for in-band managed switches: the expected uplink port (1–512) and its breakout mode. The selected port is assumed to be the first in its breakout group — i.e., the port with the SFP inserted.

- **Out of Band Management**: If enabled, denotes the switch is out of band by the management port.
- **Expected Uplink Port**: The physical port ZTP configures for in-band management (1–512). For SFP-based uplinks, set this to the port with the SFP inserted.
- **Expected Breakout**: The breakout mode for the uplink port. Consult your switch vendor documentation for supported modes.
- **Breakout Uplink Port**: Breakout uplink port identifier in the format of `1/#/#`.

Example:

- **Expected Uplink Port**: 4
- **Expected Breakout**: 4x200G
- **Breakout Uplink Port**: 1/4/3

Copyright BE Networks 2026

8.6 Working with ONTs

ONT stands for Optical Network Terminal. It is a device used in fiber-to-the-home (FTTH) or fiber-to-the-premises (FTTP) broadband setups. The ONT is located on premises and serves as the endpoint device for the optical fiber connection provided by an Internet Service Provider (ISP).

How to Create Preprovisioned Switch-Assigned ONTs from Fabrics

1. Navigate to Fabrics and select the desired fabric.
2. From the child menu, select **Switches**.
3. Select the target switch.
4. From the child menu, select **ONTs**.
5. Click **Add Preprovisioned ONT**.
6. Complete the form and click **Create Preprovisioned ONT**.

How to Create Preprovisioned Isolated ONTs from Fabrics

1. Navigate to **Fabrics** and select the desired fabric.
2. Select **Isolated ONTS**  . Click the **Add Preprovisioned ONT** button ![].
3. Complete the form that appears.
4. **Isolated ONTs** displays the preprovisioned ONT as a child. 

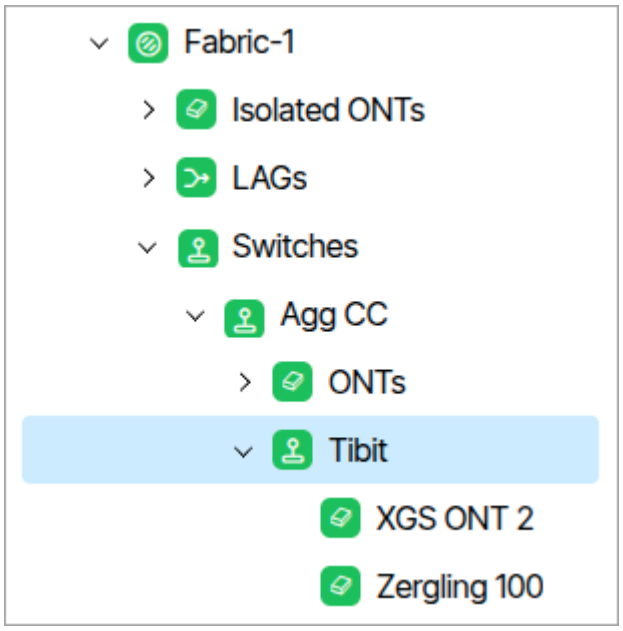
How to Create Preprovisioned ONTs From Views

1. Navigate to **Operations/Views**.
2. Click the **Add Preprovisioned ONT**.
3. Complete the form that appears.
4. The ONT now appears in any Endpoint views that are configured to show it.

8.6.1 Tibits and PONs

The system automatically recognizes these connections and provides Tibits and PONs nodes under their parent Switch nodes in the Fabrics tree.

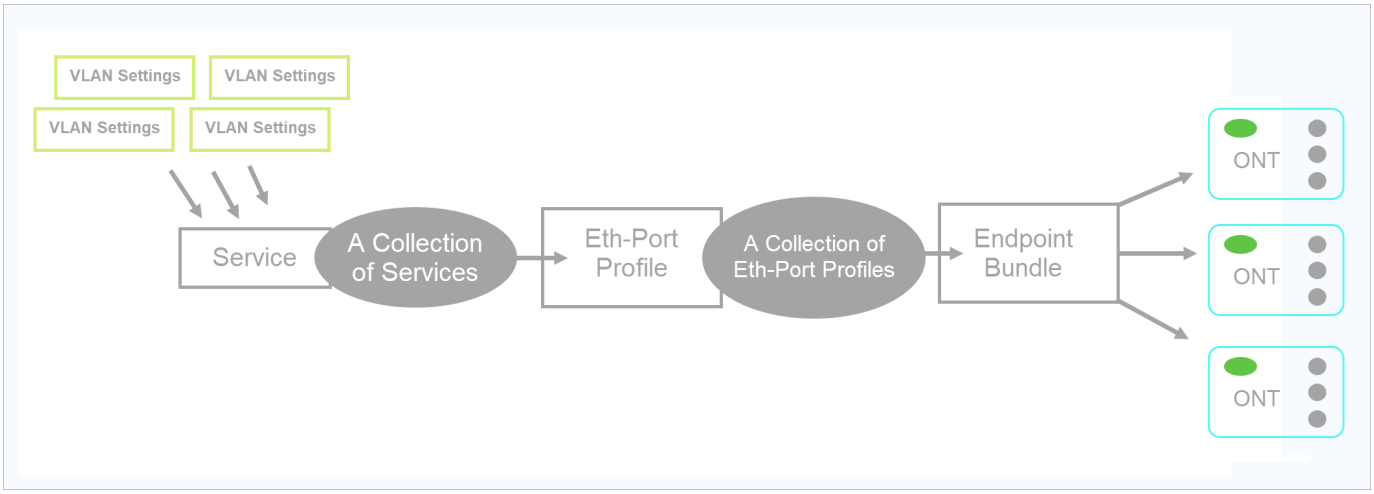
Copyright BE Networks 2026



8.6.2 How to Create a Bundle

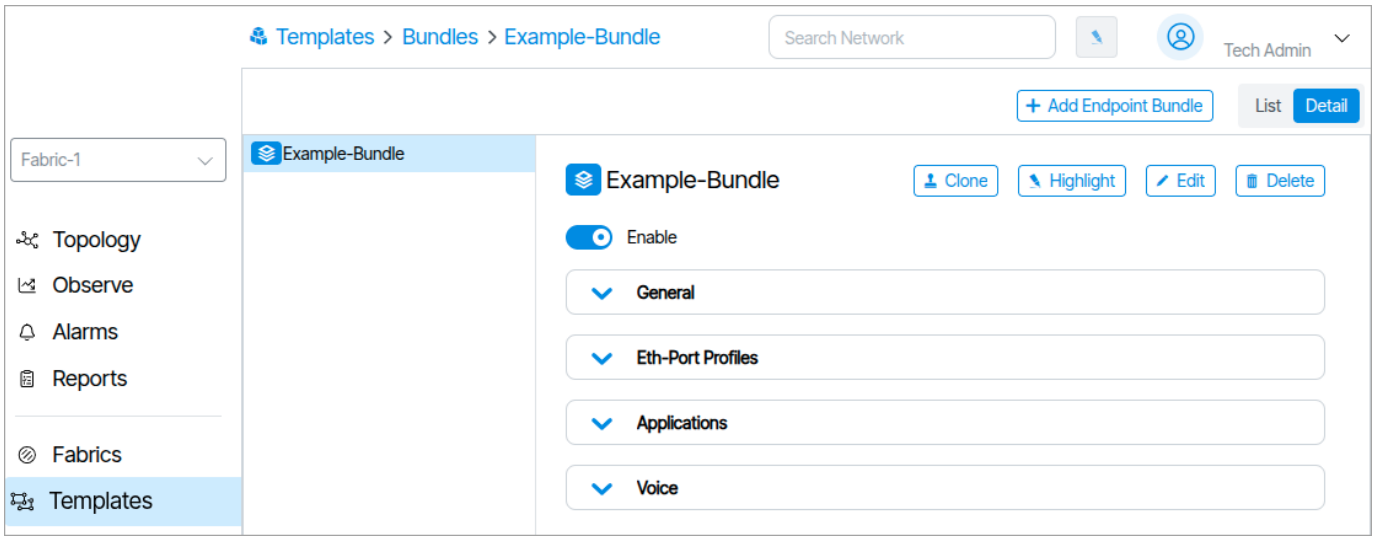
A **Bundle** is a group of settings used to configure, program and set the state of a switch or ONT. A GUI instance can have many bundles (hundreds) each for a different switch. Ideally, you will have one bundle for many ONT's. Bundle assignment is part of the provisioning process and all devices require an assigned Bundle. The following bullet points describes the relationship between settings and a **Bundle**.

- A collection of VLAN settings compose a Service.
- A collection of Services compose an Eth-Port Profile.
- A collection of Eth-Port Profiles compose a Bundle.
- A Bundle is a collection of settings transmitted to different Endpoints.



1. Go to **Templates/Bundles** and click **Add Endpoint Bundle**.
2. Type a name in the form field and click **Add Endpoint Bundle**.
3. Configure the desired settings.
4. Toggle the **Enable** switch to the "On" setting.

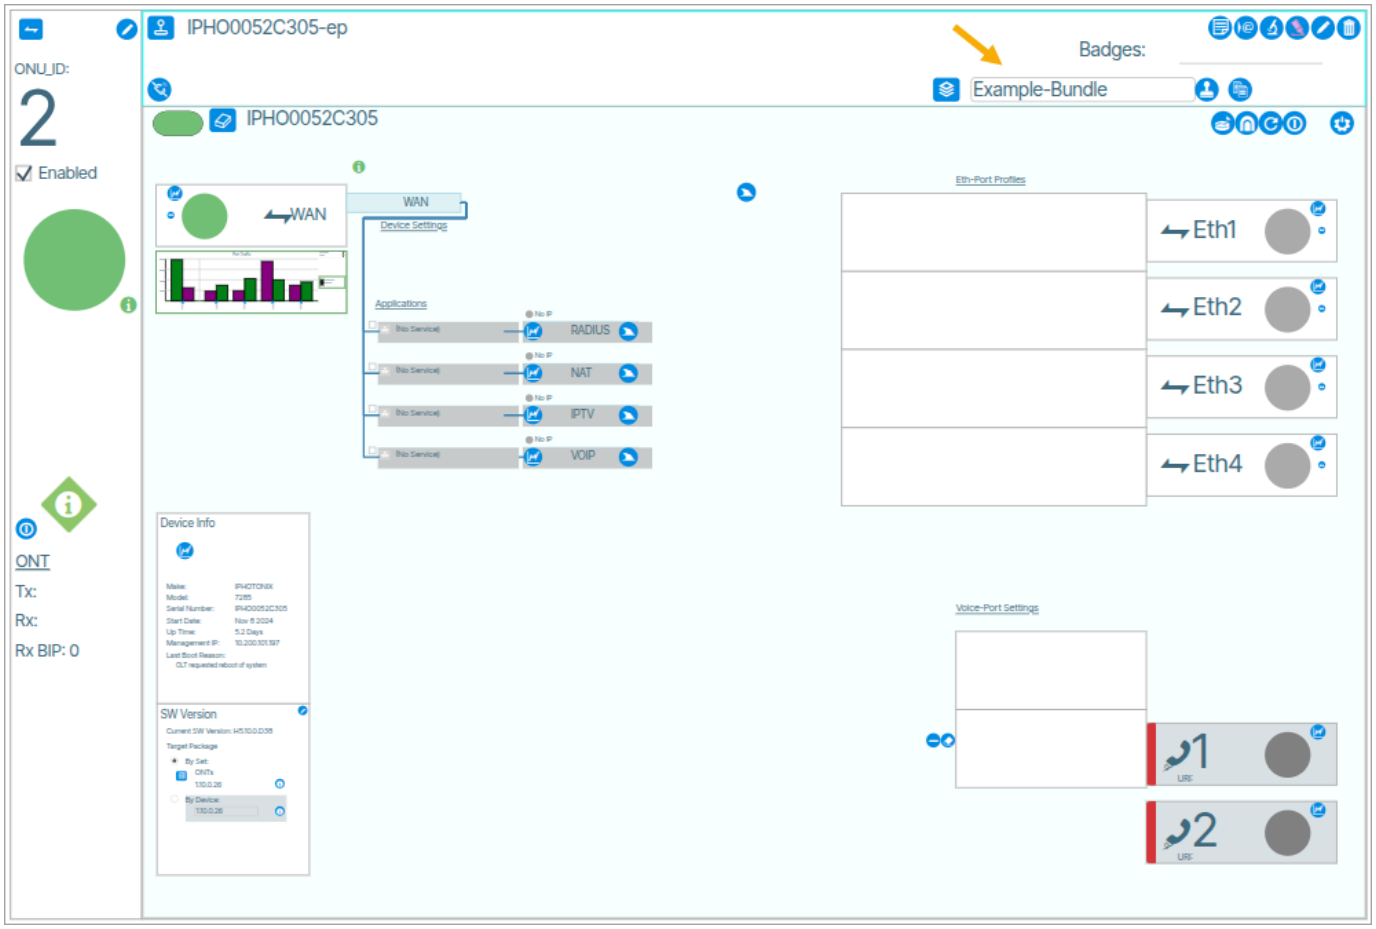
Copyright BE Networks 2026



You now assign a Bundle to an ONT.

8.6.3 How to Assign a Bundle to an ONT

To assign a Bundle to an ONT select it from the **Connected Bundle** form as shown in the following image:

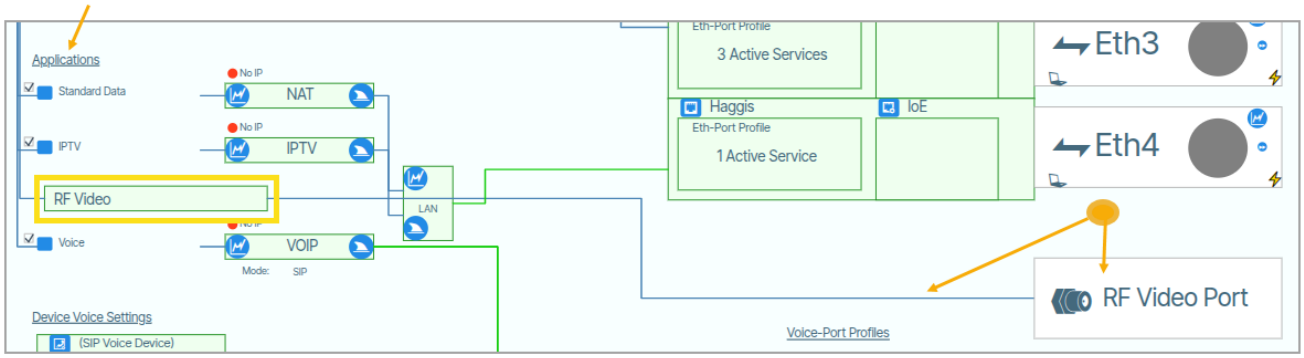


8.6.4 Drawing of Video Applications

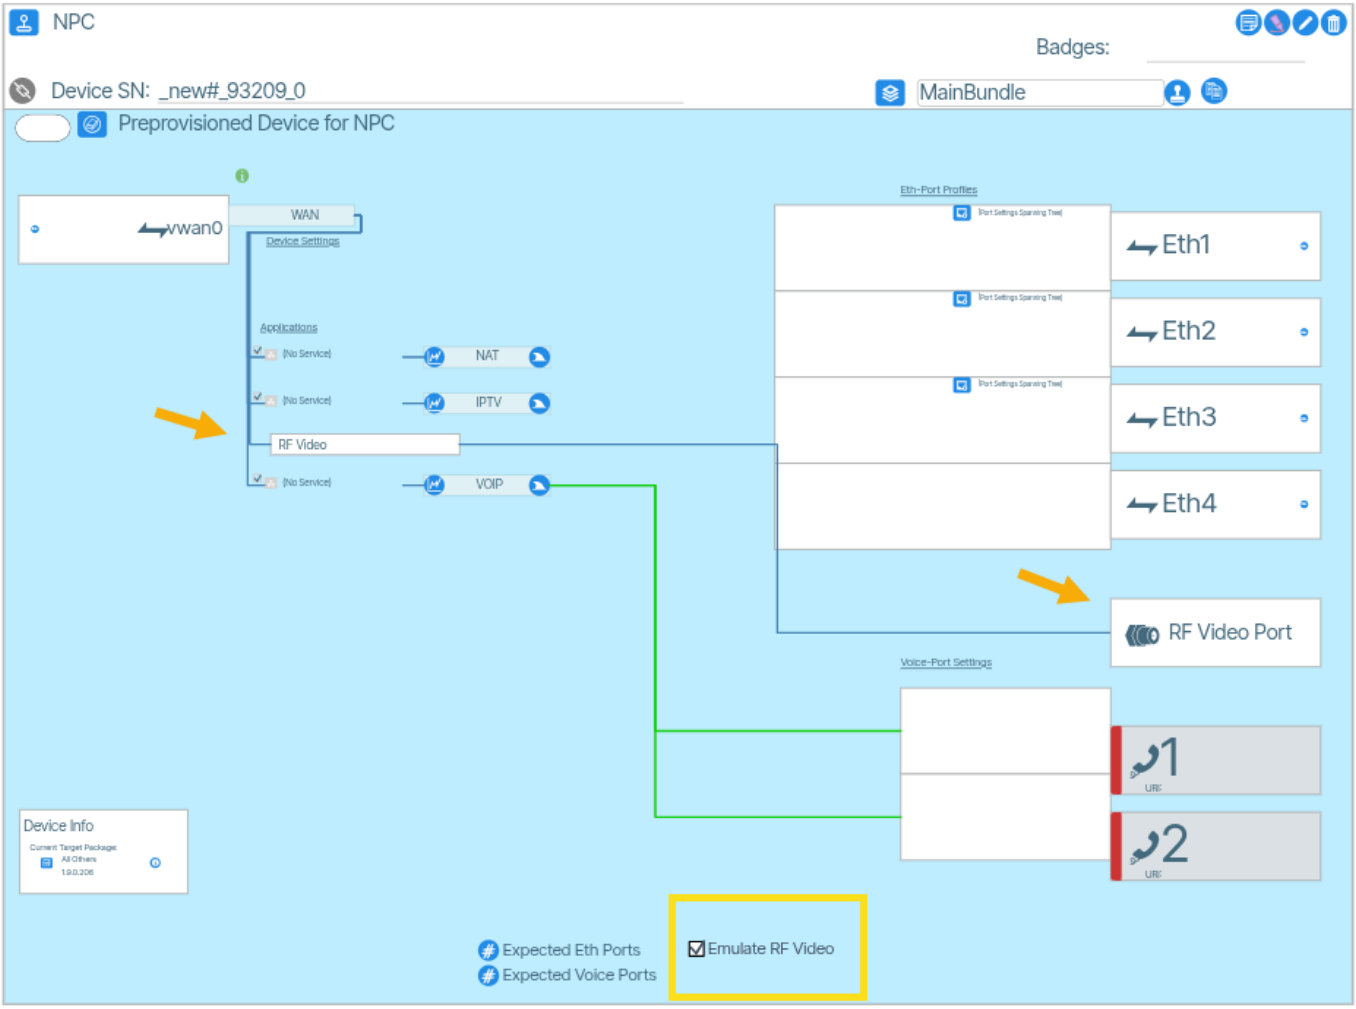
This feature displays the provisioning of the video application on a device port.

Under **Applications** there is an option for **RF Video**. A line connects that box to the **RF Video Port**. The **RF Video** field is static and cannot be edited.

Copyright BE Networks 2026



In pre-provisioned endpoints, you check the box shown below to emulate the **RF Video Port** .

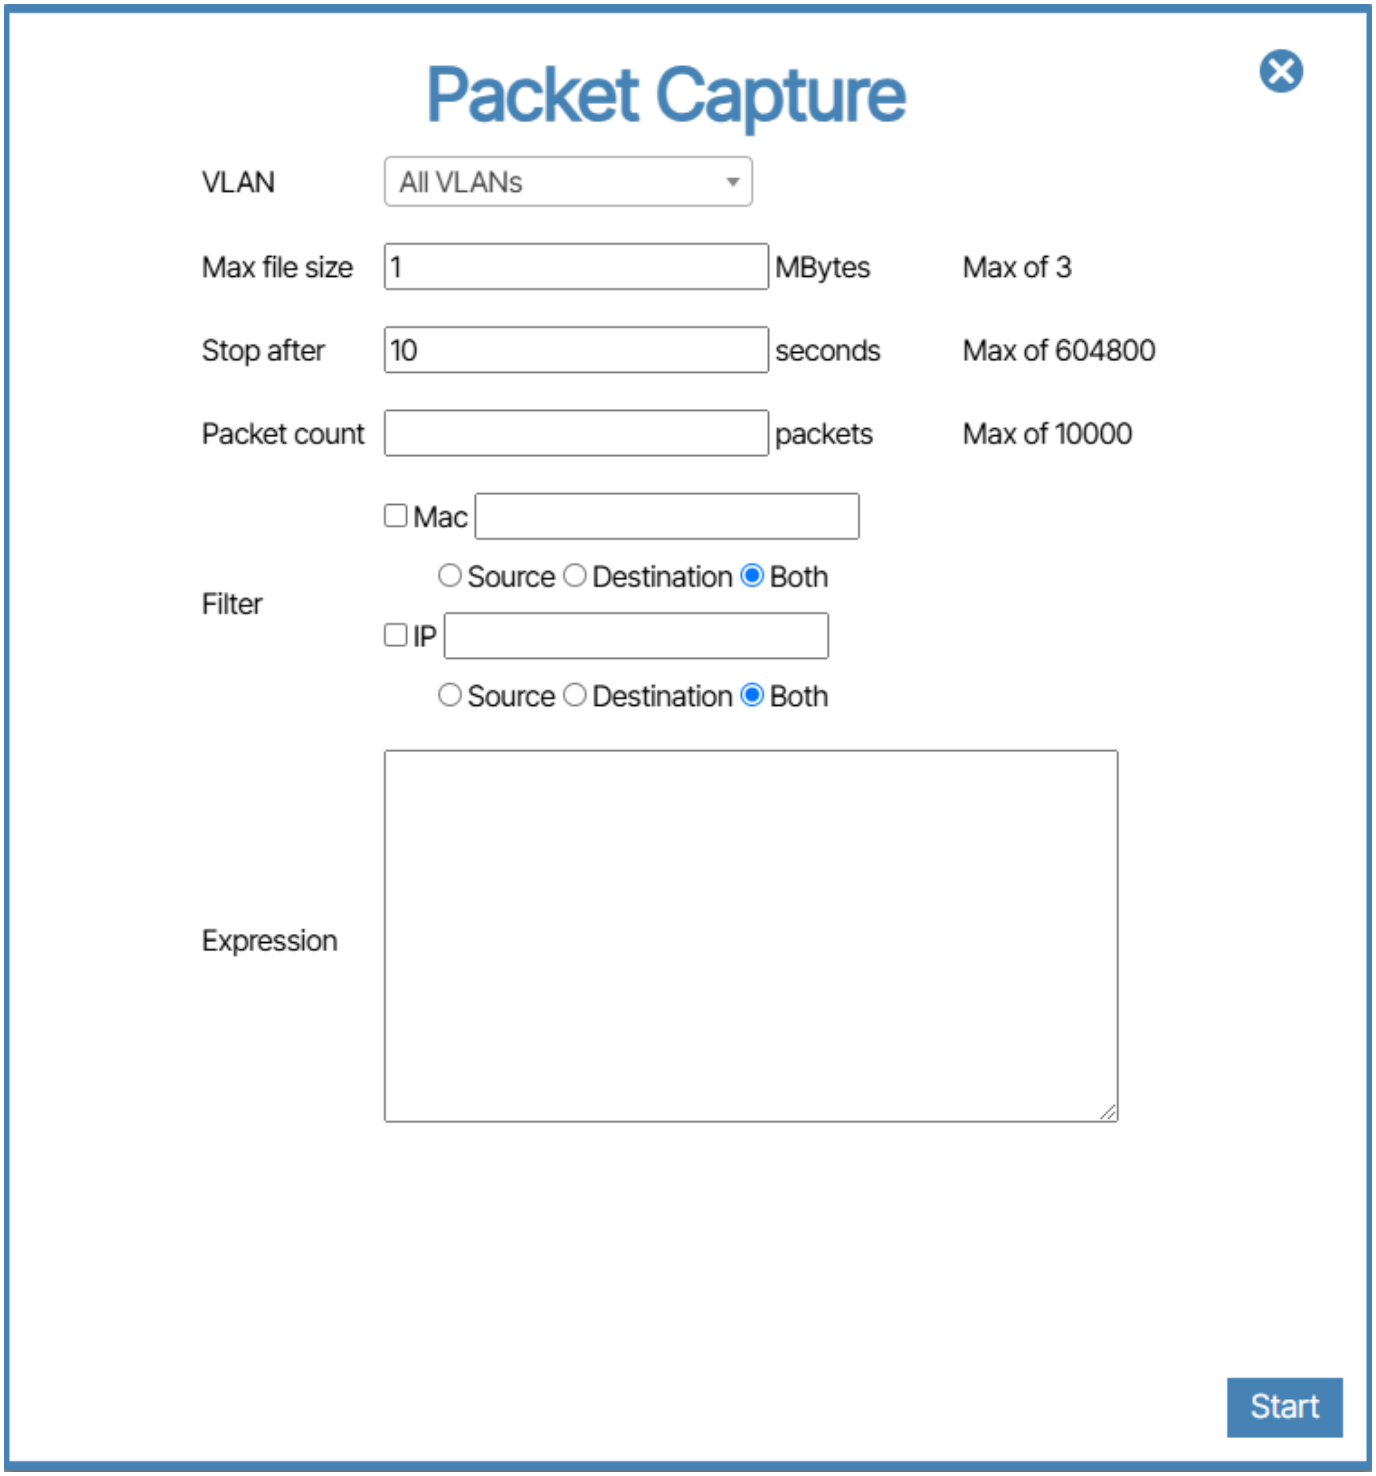


8.6.5 How to Packet Capture ONT VLAN Data

From **Topology/Network View**, zoom in on an ONT port and click the shark fin icon (). A form field will appear. To begin the packet capture fill out the field and press the **Start** button.

Copyright BE Networks 2026

Form Fields



VLAN

MAX FILE SIZE, STOP AFTER AND PACKET COUNT

When initiating a Packet Capture, the packet capture process will conclude based on the first completion among these three settings.

MAX FILE SIZE

If the file size exceeds this value before the Stop After time limit completes, or the Packet Count is reached then the packet capture stops.

Copyright BE Networks 2026

STOP AFTER

If the Max file size is not reached and the Packet Count is not reached then the packet capture will stop after the number of seconds set in this field.

PACKET COUNT

This is the number of Packets recorded. This is viewable in Wireshark. In the packet capture image below, the Packet Count is set to 10. The result is 10 packets captured.

FILTER

Mac

This field lets you capture packets of a select device based on its Mac Address. The settings Source, Destination and Either determine the packets to capture based on incoming or outgoing packets.

IP

This field lets you capture the packets of a select device based on its IP address. The settings Source, Destination and Either determine the packets to capture based on incoming or outgoing packets.

Expression

Regular Expression syntax is used here to filter select data.

8.6.6 Voice-Port Profiles

Voice-Port Profiles are *port specific* voice features and settings.

8.6.7 Device Voice Settings

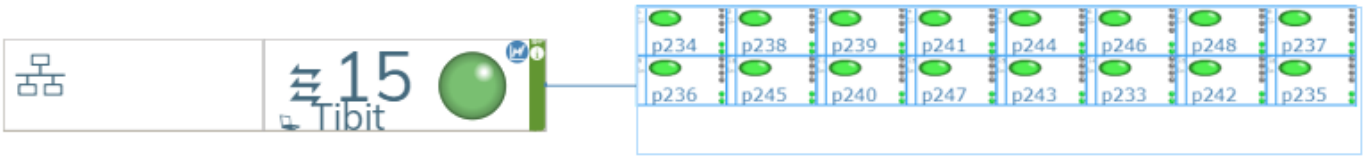
Device Voice Settings are the settings used for configuring the VOIP stack on the ONT.

8.6.8 Aggregation Switch to Endpoints

Note that three stripes in an Aggregation Switch port box indicate that the port is a “point to multi-point” connection (eg. PON).



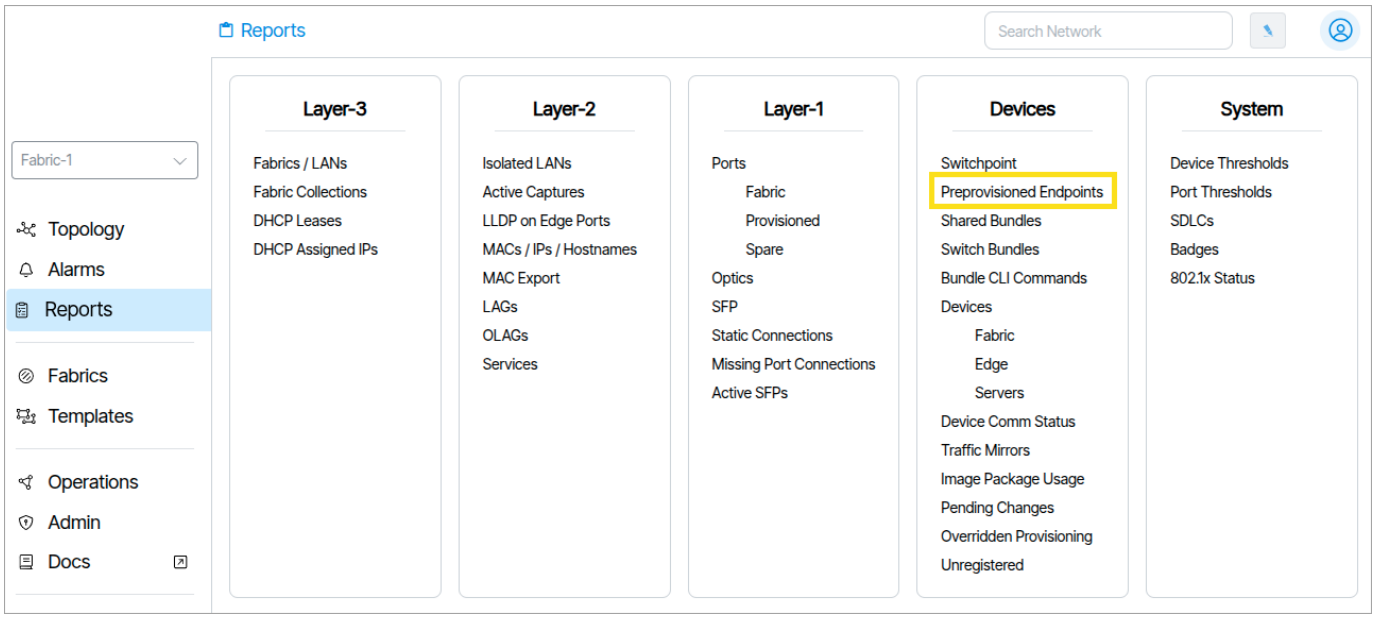
Endpoint on Aggregation Switch support point-multipoint connections



8.6.9 Preprovisioned Endpoints

Preprovisioned Endpoints is a Report that lists endpoints whose corresponding devices have not yet been identified.

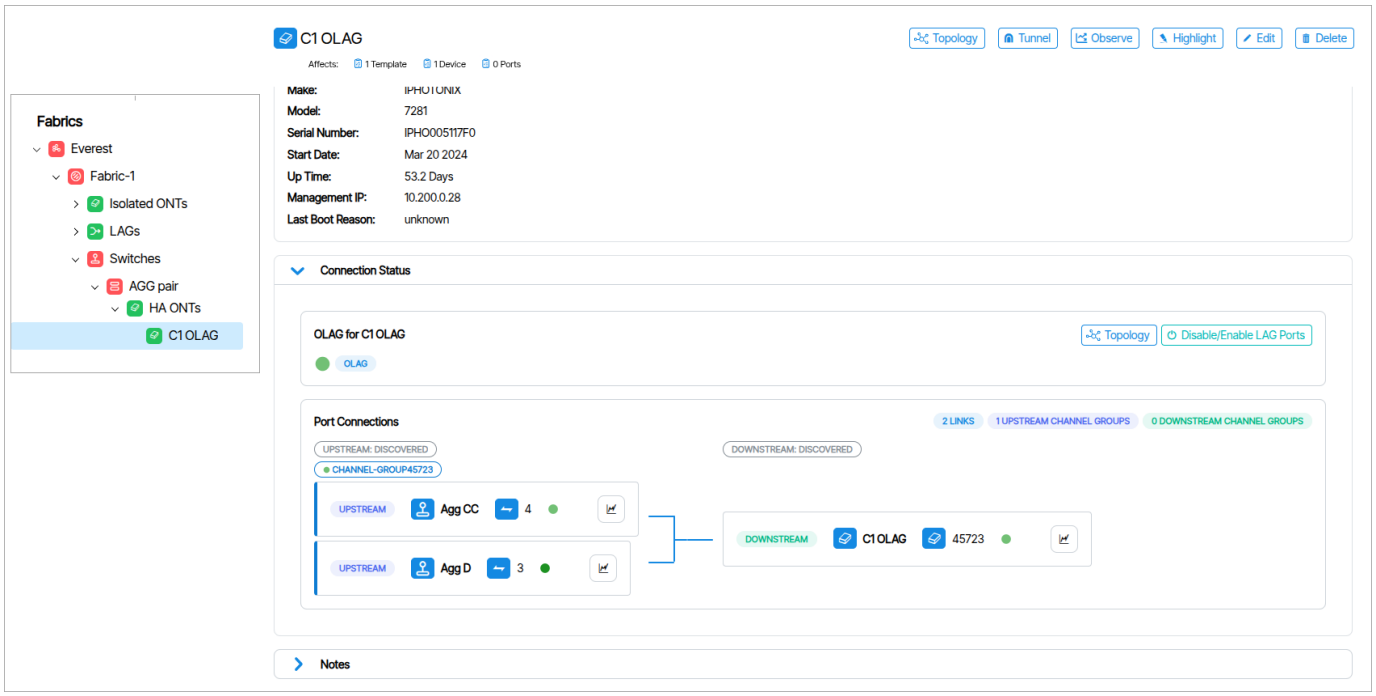
Copyright BE Networks 2026



8.6.10 Create Redundant Path to ONTs

OLAG stands for ONT-LAG and is the naming convention used for this type of redundant configuration.

To reliably initiate the OLAG feature a custom ONT vendor config file is required. Please contact customer support.



Setup Physical Hardware

To complete this part of the configuration you need:

- A pair of redundantly configured switches.
- ONT
- Joiner

Copyright BE Networks 2026

Verify Connection

In the software, make sure the linked port is marked OLAG Ready otherwise the ONT will lose communication.



Copyright BE Networks 2026

8.7 ZTP and SONiC EDGE Installation

Document Purpose

The purpose of this document is to describe the concepts and detailed instructions for deploying SONiC in Edge/Campus environments where switch management interfaces are connected using in-band connections. The process is designed to be as automated as possible while allowing for efficient and reliable installation. This document does not cover all aspects of the provisioning for the network. Please refer to the UI User Manual for the appropriate software release for this information.

Introduction

Verity for Edge is a Local Area Network (LAN) management and orchestration overlay solution that extends the power and value of Intent-Based Networking (IBN) to legacy LAN switching estates. The solution enables IT staff to focus on the high-level tasks of delivering services to endpoints – the intent – eliminating the need to focus on the time-consuming, error-prone tasks of implementation.

One of the important tasks is in installation and turn-up of new SONiC based switching devices. These devices are delivered to sites containing only the boot loader (ONIE) and require the installation of the operating system and configuration of the device to begin communications with the Verity orchestration platform. Unlike in the datacenter architecture, where management of switches is done via out of band (OOB) connections, the Edge architecture requires that switches are managed “in-band” over the same infrastructure carrying user traffic.

8.7.1 ZTP For In-Band Managed Switches in the EDGE

To successfully use the ZTP process network design decisions must be determined ahead of time. The decisions center around the connectivity expected in the fabric (single port or LAG) and the use of Multi-chassis LAG. These decisions ensure that once the switch starts the ZTP process, consistent connectivity back to the orchestration platform can be safeguarded.

> **Copper Ports Only**
>
> Please note - the fully automated process is limited to copper ports only. Switches with only SFP ports will not work in this case since setting speeds or breakout is required. For SFP only based switches used in the campus environment, ZTP must be completed with the switch set to “Managed In-band” in the Switch Endpoint and then final preparation of SFP ports and moving the connection to the in-band management VLAN must be done manually. Refer to the manufacturer’s documentation to complete the setting of in-band management on SFP ports.

Other network considerations include the management VLAN selection and tagging rules for management. In this environment, it is recommended that the management traffic is untagged on the fabric connections.

The ZTP process requires two physical connections for a switch to run the process. The first connection is on the switch’s out of band management port and the second connection is on any of the first 32 switched ports.

The ZTP process requires a DHCP server to provide addresses and DHCP options to the switches. The Verity system has an integrated DHCP server that may be used, or a DHCP in the customer network may also be used.

The ZTP process is summarized as follows. Details are shown in the subsequent sections:

1. Switch powers up and ONIE boot loader and receives **Option 114** in the DHCP response on the out of band port which points the switch to the server folder containing the appropriate code load.
2. After loading the SONiC OS, the switch reboots and again sends a DHCP request and receives **Option 67** in the DHCP response on the out of band port which points to the server folder containing the ZTP files.
3. The ZTP process populates the switch configuration based on information contained in the Device Controller managed object.
4. The ZTP process performs a connectivity test to the SDLC VM where the Device Controller is located over the connected switch port.
5. The ZTP process continues indefinitely until the connectivity test is completed successfully.
6. The Device Controller polls the network and starts communicating with the switch and then after successful discovery, the “Device Object” is created in the system.

Copyright BE Networks 2026

8.7.2 Device Onboarding Process

The steps for the process are summarized as follow:

1. Load SONIC Firmware Packages via **Admin/Software/Image Packages**
2. Set up the DHCP Server
3. Create the Device Controllers and Switchpoints
4. Bring up the switches in a serial fashion as the topology is built out.

8.7.3 Configuring DHCP

There are two options available for configuring DHCP: either use your own DHCP server or install the internal **DHCP Server Application** by running the **Add a DHCP Server** process. To use your own DHCP server, skip to the section titled **Using Your Own DHCP Server**.

The address space for the management subnet is used for the vNETC (if in same subnet), SDLC and other internal system components used to manage the devices in the system as well as the devices themselves. This subnet, that was defined during the initial VM installation process, needs a section of addresses reserved for the static assignments of the system components as well as the management interfaces of the devices in the system. The dynamically allocated range is used to support the switches through the ZTP process as well as the Device Controllers that are created to communicate with the switches.

Using the Internal DHCP Server

Double click the VNFs box  .

To setup an internal DHCP server, access the section titled **System Applications** located within the VNFs section of the Verity dashboard and click the create button  .



You are now going to install the DHCP Server. Click the **Add a DHCP Server** button  .

For the **DHCP Server**, setting the lease duration is required. The excluded address ranges are designated to reserve space for the static addresses assigned as described above. The addresses allocated for the Device Controllers and the switches during ZTP will come from within area between the excluded addresses. The subnet and this range should be sized accordingly based on how many devices are planned for the data center network. In this example, the first 100 addresses and the last 5 addresses of the management subnet have been excluded from the DHCP pool of addresses that can be allocated  .

Check the box titled **Enable** and click the checkmark to save and start the internal DHCP service  .

Using Your Own DHCP Server

To use your own DHCP server you do not need to run the **Add a DHCP Server** installer as explained in the previous section, but you need to understand the requirements described above for allocating the address space.

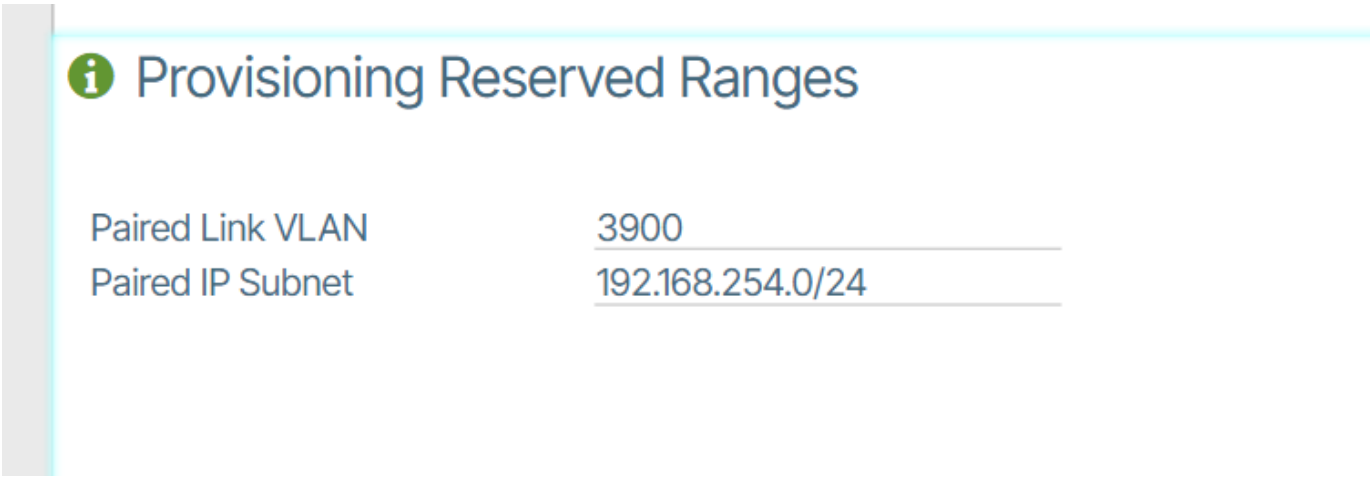
Copyright BE Networks 2026

To use your own DHCP Server the following DHCP options are required to be configured which direct the switches to the necessary files required for ZTP.

- **Option 67**: "http://\<vnetc-address or fqdn>/download/ztp/file/ztp.json"
- **Option 114**: "http://\<vnetc-address or fqdn>/download/onie/file/onie-installer"

DEFINE RESERVED RANGES

In the Admin page click **Reserved Ranges** to take you to **Provisioning Reserved Ranges**. Set up Paired Link VLAN and Paired Switch subnet (VLANs 3967-4094 are reserved in SONiC and should be avoided).



8.7.4 Load SONiC Firmware

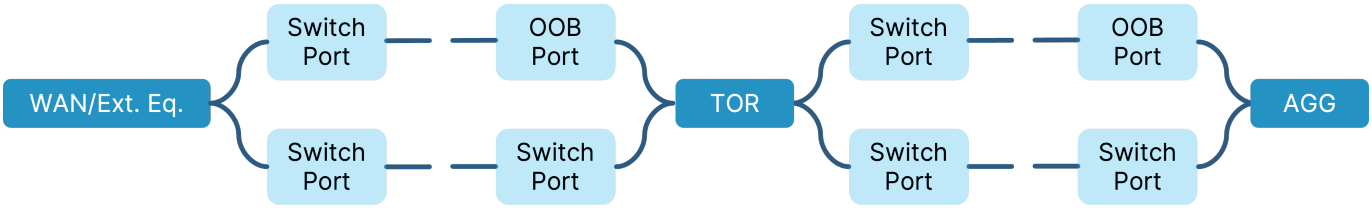
You are now going to load SONiC Firmware to VNetC. Before performing this action, you must obtain a firmware package containing the desired switch loads.

Go to the **Admin** page button located at the upper right side of the Verity dashboard UI and click **Software Packages and Licensing.** Then select **Partner Firmware Packages**. In the window that appears use the Browse Files/Drag-Drop interface to upload the pre-packaged OEM firmware file  .

After upload is complete, click the Deploy button  .

8.7.5 Building a Network without Redundancy

Edge/Campus networks are very dynamic with many topologies. For the purposes of this section, the instructions describe bringing up the TOR switch and one aggregation switch. Other aggregation switches can be connected to the TOR or to other aggregation switches, but their process is the same, if only single switches are used. Multi-chassis architectures using switch pairs is described in the next section.



The ZTP process sets the new switch up based on the contents of the Device Controller object. As a default, the first 32 ports are configured to be possible uplinks based on the management VLAN and tagging rules in the Device Controller.

ZTP requires two connections simultaneously from the switch. One of the first 32 switch ports is designated to be the uplink connection to the WAN/ Orchestration platform. The “out of band port” or “Management Port” of the switch is used to handle ONIE download and the ZTP process. Accommodations must be provided to allow the first switch (TOR) to go through the ZTP process. Once the TOR is up and running, connections for subsequent switches are made through the TOR. It is recommended that the management connections in the switching fabric are all untagged.

Copyright BE Networks 2026

TOR

BUILD THE DEVICE CONTROLLER:

1. Enter LLDP string representing the Service Tag/Serial Number of the switch.
2. Enter the remaining fields as shown in this example:  .
3. Turn the device controller on (Power switch ICON)
4. Connect one of the first 32 switch ports to an untagged upstream port with a path/route to the orchestration platform.
5. Connect “out of band” port of the switch to an untagged upstream port with a path/route to the orchestration platform.
6. Upon completion of the ONIE boot load and ZTP process, the Device controller will locate the switch and it will draw on the system UI.
7. Remove the wire connected on the out of band management port of the switch.

AGG

BUILD THE DEVICE CONTROLLER:

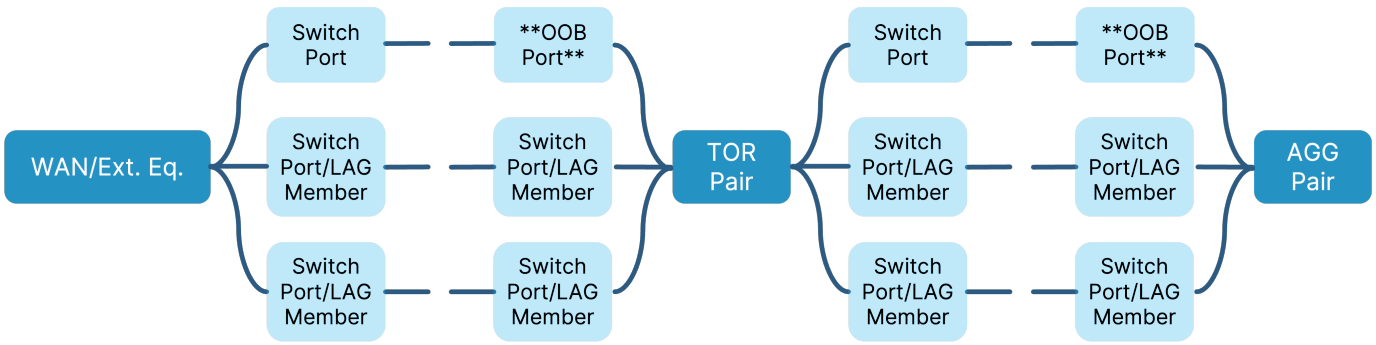
1. Enter LLDP string representing the Service Tag/Serial Number of the switch.
2. Enter the remaining fields as shown in this example:  .
3. Create an Ethernet Port Profile with the site’s management VLAN in it and place it on a selected switch port on the TOR.
4. Turn the device controller on (Power switch ICON)
5. Connect one of the first 32 switch ports of the AGG to the TOR port with the management Ethernet Port Profile from step 3).
6. Connect “out of band” port of the switch to an untagged upstream port with a path/route to the orchestration platform.
7. Upon completion of the ONIE boot load and ZTP process, the Device controller will locate the switch and it will draw on the system UI.
8. Remove the wire connected on the out of band management port of the switch.

Repeat the AGG steps for other switches connected to the TOR, AGG and any other switches in the site topology. The management profile can be placed on any ports in the switches already discovered as the network infrastructure is built up.

8.7.6 Building a Multi-Tier Multi-chassis-LAG network

Building a multi-chassis LAG network has inherent complexities to maintain connectivity to all devices as the network build out progresses. Decisions regarding the switch pairs must be pre-planned so that the Port Channel assignments created during ZTP can be maintained throughout the build out. Additionally, careful attention **MUST** be given to the order of the pairing of switches and creation of uplinks to both maintain connectivity as well as prevent traffic loops. There are two steps where the A side of the pair must be placed in read only mode to ensure that the configuration gets to the B side first.

For the purposes of this discussion, paired switches will have an “A” side identifying the first of the switch pair to be brought into the infrastructure, and the “B” side identifying its partner switch.



Copyright BE Networks 2026

The general process is similar to the non-redundant switch bring-up with respect to ZTP. The following issues highlight the process for the redundant fabric.

1. The ZTP process creates Port Channel groups including the first 32 switch ports so that any of those ports can be used to initiate the network build out.
2. A sides are brought up first.
3. B sides are brought up as “child” switches to their respective A side.
4. After B is brought up, A and B are paired.
5. After pairing is complete, uplinks from the B side to the parent’s pair are added.
6. LAG objects are created when the Device Controllers are created.
7. LAG objects for connections between Parent and Child switches are manually provisioned on the parent side and automatically assigned by the Verity software on the Child side.
8. Intra pair connections are initially Parent-Child and then become manually provisioned on both sides during the pairing process.
9. Four physical connections are expected between switch pairs (A-A, A-B, B-A and B-B)
10. Intra switch LAG should have two physical links at a minimum.

Detailed Process Steps:

TOR-A

BUILD THE DEVICE CONTROLLER:

1. Enter LLDP string representing the Service Tag/Serial Number of the switch.
2. Enter the remaining fields as shown in this example:  . Take note that the TOR-A is designated as connected via LAG by checking the option.
3. TOR-A create a new uplink LAG, with and uplink service port profile, and assign to one of the first 32 ports of the pre-provisioned TOR-A Switchpoint. Please note that the LAG indication is a checkbox along with the “Located By” setting.
4. Connect the port assigned in the Switchpoint TOR-A to upstream device with Port Channel already configured with the Management VLAN with connectivity to Verity orchestration platform.
5. Connect TOR-A out of band management port to an untagged upstream port for ONIE/ZTP processing on the same Management VLAN.
6. Power on the Device Controller and wait for the completion of the ZTP process and the switch to be discovered by the system. The default behavior for new switches is Read Only mode. Once the switch is discovered, disable read only mode.

TOR-B

BUILD THE DEVICE CONTROLLER:

1. Enter LLDP string representing the Service Tag/Serial Number of the switch.
2. Enter the remaining fields as shown in this example:  .
3. Note that the switch is located by “Peer LAG.” The type of LAG is referred to as an “IDL” LAG
4. Assign the LAG created in the TOR-B Device Controller to a port on TOR-A
5. Connect the link where the LAG was created on TOR-A to any one of the first 32 ports of TOR-B. (Note: This connection will eventually become the intra chassis connection between A and B sides once the switches are paired.
6. Create an Ethernet Port Profile with the site’s management VLAN in it, untagged, and place it on a selected switch port on the TOR-A.
7. Connect the out of band management port of TOR-B to the port for ONIE/ZTP processes.
8. Power up the switch and the Device controller and wait for the ZTP process to complete and the switch to be discovered by the system. Disable read only mode for the TOR-B switch.

PAIR TOR-A and TOR-B

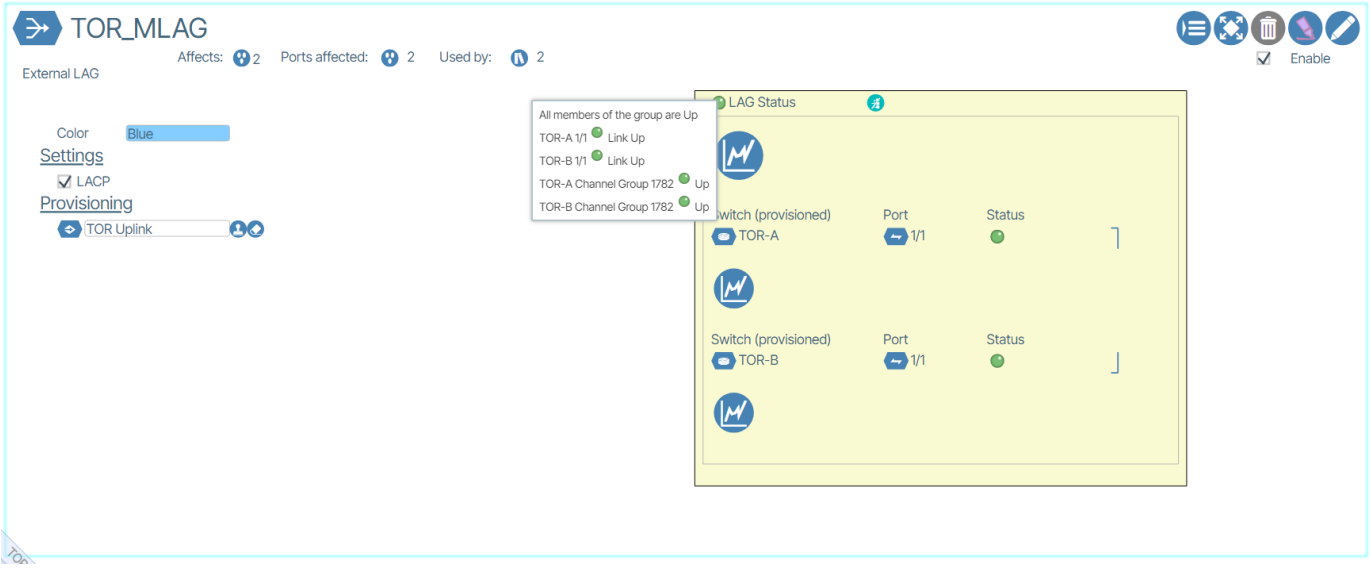
1. Remove out of band connection link from TOR-B

Copyright BE Networks 2026

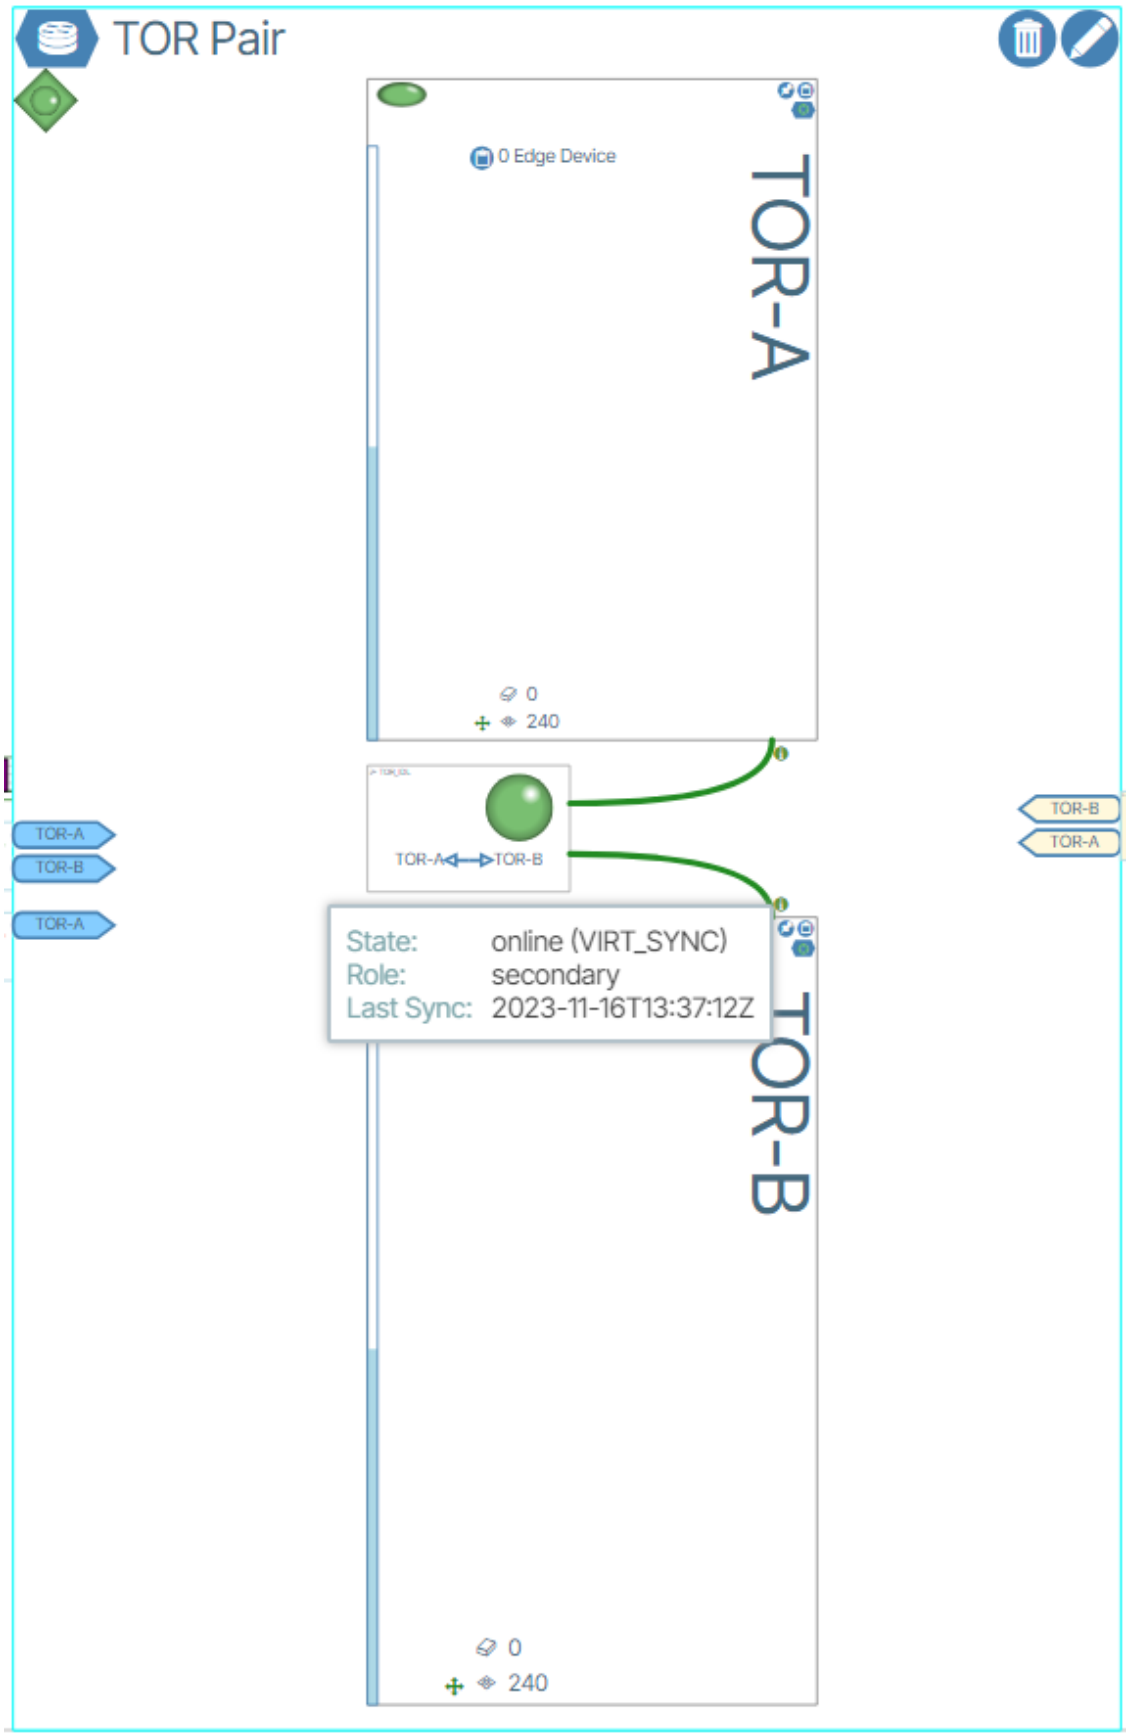
2. **IMPORTANT: Place TOR-A in READ ONLY mode.**
3. Pair TOR-A/B using the dialogue box in Network View  .
4. Assign 2nd IDL port on TOR-A to the IDL LAG created in the TOR-B controller and assign two links on TOR-B. The first is the link that was used for the first connection, and the second can be any other port on TOR-B.
5. Wait for TOR-B to complete provisioning (light green back to white) to ensure that the configuration is on the B side.
6. **Place TOR-A back into read/write mode and wait for it to complete provisioning**
7. Wait for TORs to get into synchronization with each other as identified by the green “i” dots drawn next to the connections between them.

Add the remaining Uplinks to the TOR pair

1. **IMPORTANT: Place TOR-A in READ ONLY mode.**
2. Assign TOR_MLAG uplink to any available port on TOR-B and wait for provisioning to complete.
3. Connect uplink from TOR-B to upstream device with the same Port Channel connected to TOR-A
4. Wait for the system to provision the LAG connection into the B side. It is possible that A or B may lose communications with the system once the LAGs are converted to multi-chassis LAGs.
5. **Clear the READ ONLY mode on TOR-A. Wait for the system to stabilize and all communications are up for both switches.**
6. LAG status should show all links up and channel group up within a minute and IDL LAG should show up and paired switches sync.



Copyright BE Networks 2026



1. More links can be added to TOR-A and TOR-B to make cross connections to the device above the TOR pair for full switch and link redundancy.

Copyright BE Networks 2026

8.7.7 Multi-Chassis Aggregation Switches -A/B pair

AGG-A

BUILD THE DEVICE CONTROLLER:

1. Enter LLDP string representing the Service Tag/Serial Number of the switch.
2. Enter the remaining fields as shown below in this example:  .
3. Add the LAG created in the device controller to a port on TOR-A.
4. Connect uplink from AG-A on any of the first 32 ports to TOR-A port where the LAG is provisioned.
5. Connect out of band management port of AG-A to an untagged upstream port for ZTP using the Ethernet Port Profile created in the previous section.
6. Power up Device Controller
7. Wait for successful ZTP indicated by the device getting discovered by the system.

AG-B

BUILD THE DEVICE CONTROLLER:

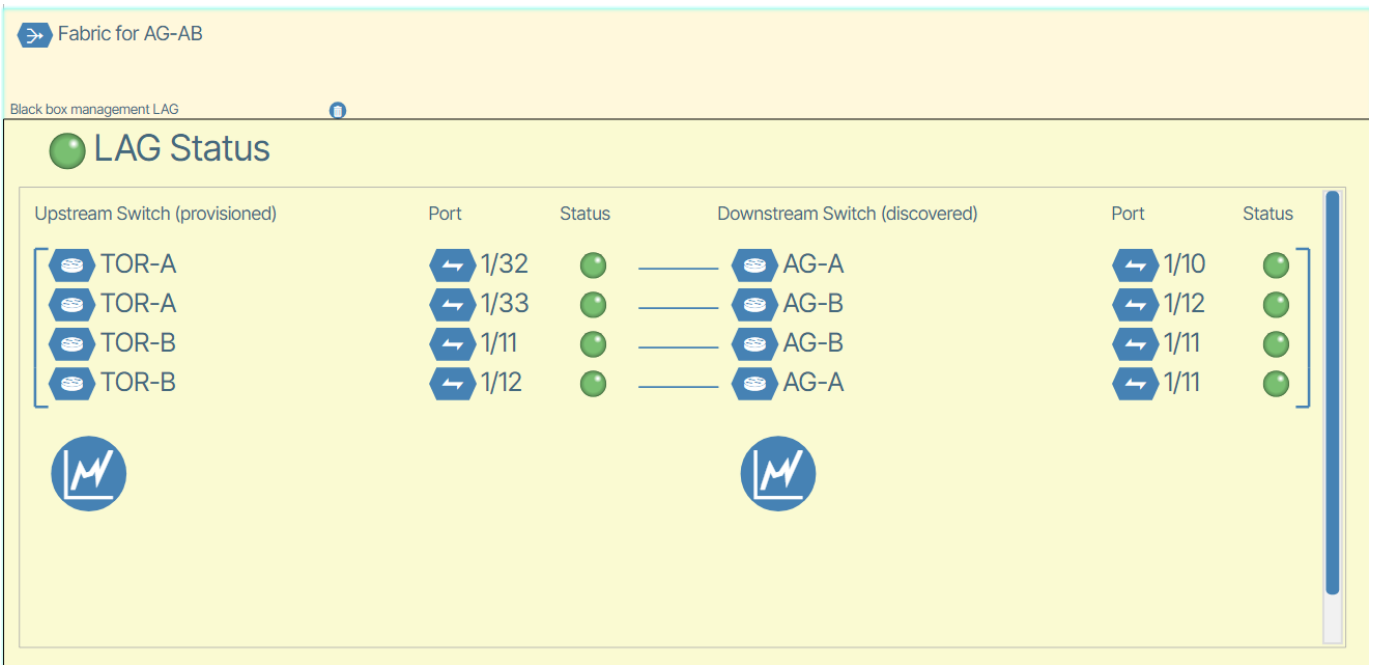
1. Enter LLDP string representing the Service Tag/Serial Number of the switch.
2. Enter the remaining fields as shown below in the following example:  .
3. Note that the switch is located by “Peer LAG.” The type of LAG is referred to as an “IDL” LAG
4. Assign the LAG created in the AG-B Device Controller to a port on AG-A
5. Connect the link where the LAG was created on AG-A to any one of the first 32 ports of AG-B. (Note: This connection will eventually become the intra chassis connection between A and B sides once the switches are paired.
6. Create an Ethernet Port Profile with the site’s management VLAN in it, untagged, and place it on a selected switch port on any in service switch that can be reached.
7. Connect the out of band management port of AG-B to the port for ONIE/ZTP processes.
8. Power up the switch and the Device controller and wait for the ZTP process to complete and the switch to be discovered by the system.

8.7.8 PAIR AG-A and AG-B

1. Remove out of band connection link from AG-B
2. **IMPORTANT: Place AGG-A in READ ONLY mode.**
3. Pair Agg-A/B using the dialogue box in Network View  .
4. Assign 2nd IDL port on AG-A to the IDL LAG created in the AG-B controller and assign two links on AG-B. The first is the link that was used for the first connection, and the second can be any other port on AG-B.
5. Wait for AG-B to complete provisioning (light green back to white) to ensure that the configuration is on the B side.
6. **Place AG-A back into read/write mode and wait for it to complete provisioning**
7. Wait for AGs to get into synchronization with each other as identified by the green “i” dots drawn next to the connections between them.
8. **IMPORTANT: Place AG-A in READ ONLY mode!**
9. Assign the LAG created in the AG-A Device Controller to any available port on TOR-B and wait for provisioning to be completed.
10. Connect uplink from AG-B to TOR-B’s provisioned port
11. Wait until provisioning to be completed on AG-B. There may be cases where AG-B temporarily loses communications with orchestration platform. If this occurs, wait one minute.
12. Clear the READ ONLY mode on AG-A.
13. Additional uplinks should be added between AG-B and TOR-A and AG-A and TOR-B. Add the LAG manually on the TOR side, connect the wires, and the system automatically adds the new links to the LAG on the AG side.

Copyright BE Networks 2026

14. LAG status should show all links up and channel group up within a minute and IDL LAG should show up and paired switches sync.



More tiers may be added below AGG AB pair and other AGG pairs may be added to the same tier. Follow the same procedure as described.

8.7.9 Using a High-Speed or Breakout Port as the Uplink

By default, ZTP provisions the first 32 switch ports with the management VLAN, so a copper uplink on any of those ports needs no additional settings. However, most campus installations use higher-speed ports for uplinks, and some use switches without copper ports at the TOR (core/aggregation) layer. For these designs, you specify the uplink port and its speed or breakout in the switch's Network Settings.

You must set the uplink port settings when either of the following is true:

- The uplink is on an SFP-based (non-copper) port.
- The uplink is on a port above port 32.

Switches without copper ports fall under the first condition. With these settings, they complete ZTP automatically, without the manual procedure in the Copper Ports Only note at the beginning of this document.

Uplink Port Settings

These settings are in the **Device Management** section of the switch.

Setting	Description
Out of Band Management	If enabled, the switch is managed out of band through its management port. Leave this setting disabled for in-band managed switches.
Expected Uplink Port/Port Group	The port number of the uplink (1–512). On switches whose SFP ports are organized in port groups, enter the port group number instead.
Expected Breakout/Port Group Speed	The breakout mode of the uplink port (for example, `1x10G`), or the speed of its port group (for example, `10G`). Select **None** if the port uses neither. Supported modes vary by switch; refer to the switch vendor's documentation.
Uplink Ethernet Port	The interface the uplink cable connects to: `1/<port>` for a port without a breakout, or `1/<port>/<lane>` for a breakout lane.

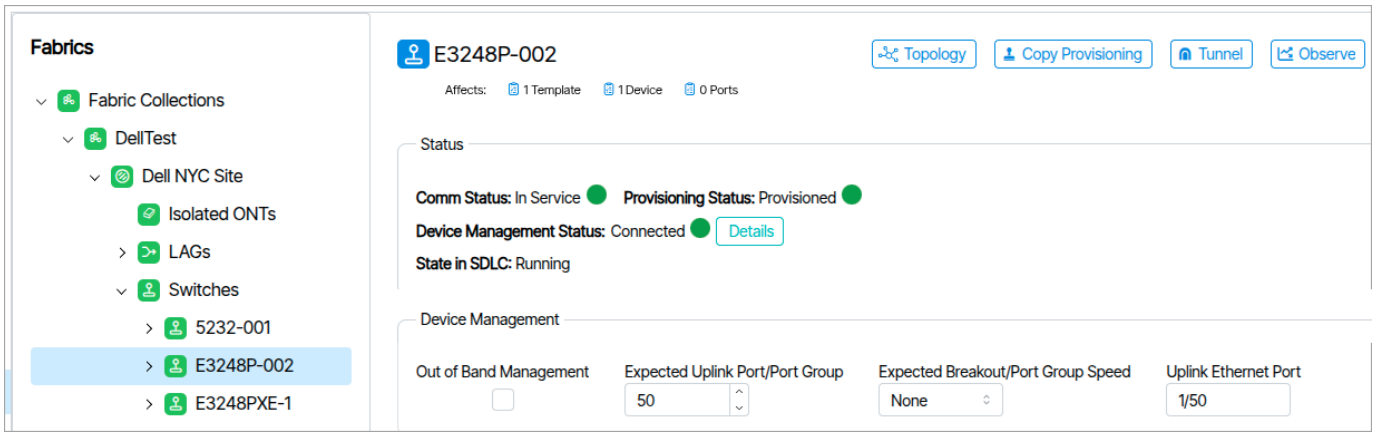
Copyright BE Networks 2026

How you fill in these settings depends on the type of uplink port:

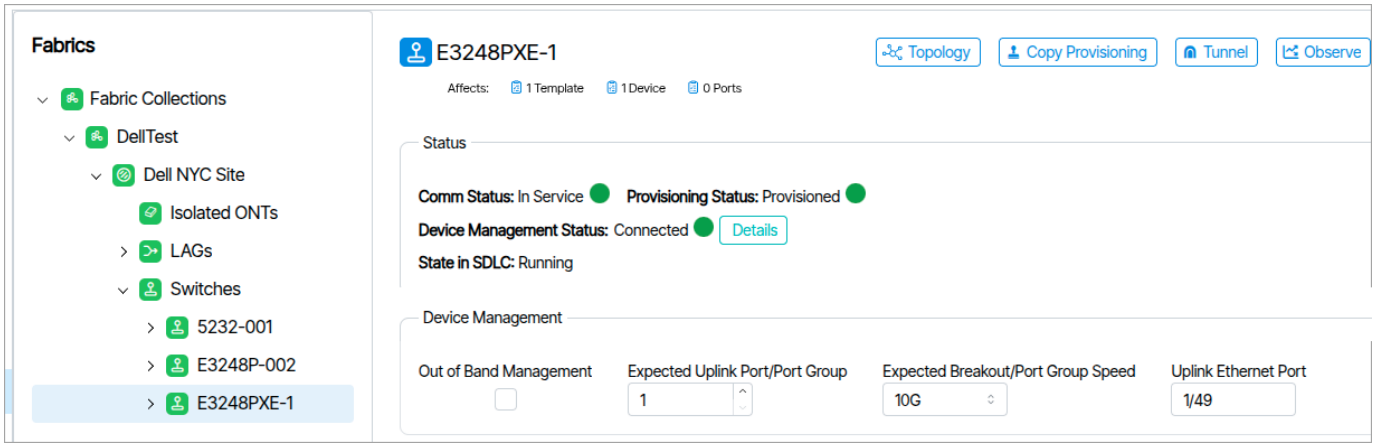
Uplink type	Expected Uplink Port/ Port Group	Expected Breakout/Port Group Speed	Uplink Ethernet Port
Plain port (for example, port 50 on Device E3248P-002)	Port number: `50`	`None`	`1/50`
Port in a port group (for example, port 49 in port group 1 on Device E3248P-002)	Port group number: `1`	Port group speed: `10G`	`1/49`
Breakout port (for example, port 1 on Device 5232F-001)	Port number: `1`	Breakout mode: `1x10G`	`1/1/1`

Plug the SFP and uplink cable into the interface named in **Uplink Ethernet Port**.

Plain port

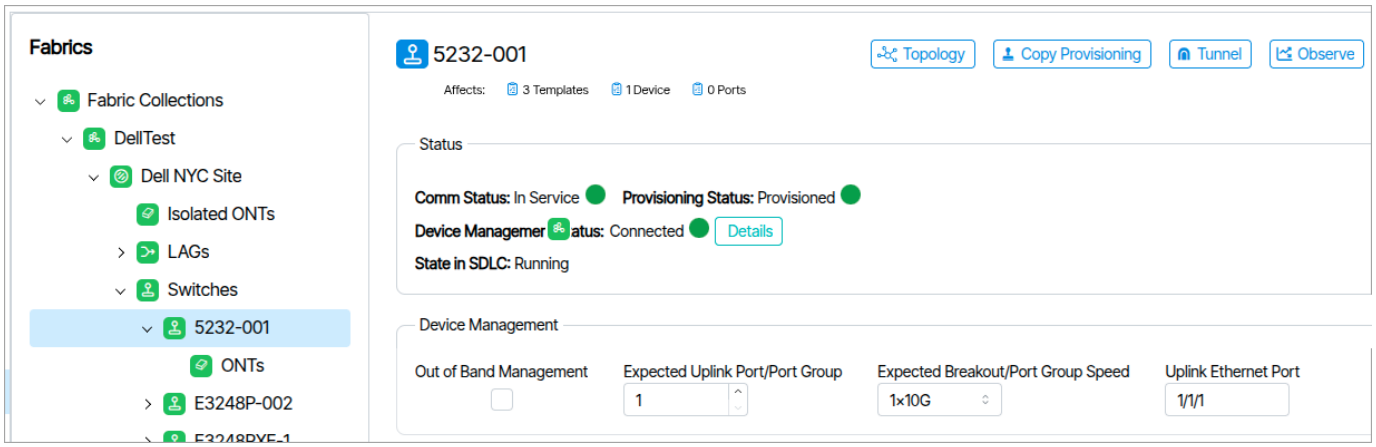


Port in a port group



Breakout port

Copyright BE Networks 2026



How It Works

When the Expected Uplink Port/Port Group is set, ZTP configures only that port for management. ZTP adds the port with the breakout you specified and applies the management VLAN to that port alone. The first 32 ports are **not** provisioned with the management VLAN, so the uplink works only when it is connected to the port you specified.

ZTP still runs over the out of band management port, so the out of band connection is required as usual. After SONiC is installed and the switch configuration is applied, ZTP tests connectivity to the Device over the specified uplink port. After the test succeeds, Verity manages the switch over that port.

Bring Up the Switch

1. On the Switchpoint for the switch, make sure **Out of Band Management** is disabled, then set **Expected Uplink Port/Port Group** and, if the port uses a breakout, **Expected Breakout/Port Group Speed** and **Uplink Ethernet Port**.
2. Follow the TOR or AGG procedure in Building a Network without Redundancy, with one change: wherever a step says to connect one of the first 32 switch ports, connect the expected uplink port instead. All other steps, including connecting the out of band management port, stay the same.

Verify the Uplink and Remove the Out of Band Cable

No indicator shows when it is safe to remove the out of band cable. Use this check instead:

1. Wait until the switch is discovered and drawn in the system UI and is stable.
2. Remove the cable from the out of band management port.
3. Confirm that the switch stays online.

If the switch goes offline after you remove the cable, in-band management is not configured correctly. Check that the uplink cable is in the Expected Uplink Port/Port Group, that the Expected Breakout/Port Group Speed matches the cable or optic, and that the upstream port carries the management VLAN. If the problem continues, contact support.

> **Do Not Remove the Management VLAN from the Uplink**
>
> The management VLAN on the uplink port is how Verity communicates with the switch. If you delete the service interface or port profile that carries the management VLAN on the uplink, the switch loses its management connection. To recover, you must connect to the switch locally and reconfigure the uplink port with the

Copyright BE Networks 2026

9. Getting Help - Report a Problem

If you find any problems or need additional assistance with this documentation, please feel free to reach us by filling out the form below.

Thank you!

Copyright BE Networks 2026