



## **Fabric Services System**

# **USER GUIDE RELEASE 22.4**

**3HE 18705 AAAA TQZZA**

**Issue 1.0**

**May 2022**

**© 2022 Nokia.**

Use subject to Terms available at: [www.nokia.com/terms/](http://www.nokia.com/terms/).

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

---

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

© 2022 Nokia.

# Table of contents

<b>1 About this document.....</b>	<b>11</b>
1.1 What's new.....	11
1.2 Precautionary and information messages.....	11
1.3 Conventions.....	11
<b>2 Overview.....</b>	<b>13</b>
2.1 Concepts.....	13
2.2 Life cycle of a fabric.....	14
2.2.1 Design the fabric intent.....	14
2.2.2 Model in the Digital Sandbox.....	15
2.2.3 Deploy the fabric intent.....	16
2.2.4 Monitor the fabric intent.....	16
2.2.5 Repeat the cycle.....	17
2.3 Workload VPN intents.....	17
2.4 Maintenance intents.....	17
2.5 Additional resources.....	18
2.5.1 Documentation resources.....	18
2.5.2 Accessing the API descriptions from the system GUI.....	18
<b>3 User interface basics.....</b>	<b>20</b>
3.1 Signing in to the Fabric Services System.....	20
3.2 Signing out of the Fabric Services System.....	20
3.3 The dashboard.....	20
3.3.1 Notifications.....	24
3.4 The main menu.....	25
3.5 Detail panels.....	27
3.6 Lists.....	28
3.6.1 Custom filters.....	30
<b>4 Deployment regions.....</b>	<b>32</b>
4.1 Region properties.....	32
4.2 Creating a region.....	34
4.3 Modifying a region.....	35
4.3.1 Implications of modifying a region.....	36

4.4 Updating a region to use iBGP.....	37
4.4.1 Updating region settings to use iBGP.....	37
4.4.2 Configuring fabric nodes as route reflectors.....	38
4.4.3 Updating existing fabric intents to use new iBGP settings.....	38
4.5 The region map.....	39
4.5.1 Viewing and using the region map.....	40
4.5.2 Viewing and using the region list.....	40
4.5.3 Manipulating the region map.....	41
4.6 The deployment pipeline.....	41
4.6.1 Viewing the deployment pipeline.....	44
4.6.2 Working with the deployment pipeline.....	44
4.7 The Fabric Intent List view.....	45
4.8 The Topology Map view.....	46
<b>5 Fabric intents.....</b>	<b>48</b>
5.1 Nokia certified templates.....	48
5.2 Fabric topology.....	48
5.2.1 Flexible leaf/spine fabrics.....	49
5.2.1.1 The Nokia certified flexible leaf/spine template.....	49
5.2.2 Backbone fabrics.....	50
5.2.2.1 Border leaf.....	51
5.2.2.2 The Nokia certified backbone template.....	51
5.3 The design-and-deploy workflow.....	51
5.3.1 Create the fabric intent.....	52
5.3.2 Explore the fabric intent.....	53
5.3.3 Create and assign labels.....	53
5.3.4 Associate nodes with hardware.....	53
5.3.5 Add the fabric intent to the deployment pipeline.....	53
5.3.6 Deploy the fabric intent to hardware.....	53
5.4 Supported hardware.....	53
5.5 Software and image catalogs.....	54
5.6 Fabric intents page.....	54
5.7 Creating a new fabric intent: flexible leaf/spine.....	56
5.8 Adding a second type of hardware for leaf nodes.....	59
5.9 Creating a new fabric intent: backbone.....	60
5.10 Viewing a fabric intent.....	63

5.10.1 Groups.....	63
5.10.2 Information displays.....	65
5.10.3 Error indicators.....	66
5.10.4 Viewing the event log.....	67
5.10.5 Display controls.....	68
5.10.6 Viewing a fabric intent as code.....	69
5.10.7 Viewing the configuration file for a single node.....	70
5.10.8 Downloading the initial node configuration.....	71
5.10.9 Viewing the fabric inventory.....	71
5.10.10 Viewing fabric links.....	72
5.10.11 Viewing edge links.....	72
5.10.12 Generating a wiring plan.....	73
5.11 Creating and managing LAGs.....	73
5.11.1 Creating LAGs.....	73
5.12 Configuring a breakout port.....	75
5.13 Modifying a fabric intent.....	76
5.13.1 Editing a fabric intent.....	76
5.13.1.1 Effect of edits on a related backbone intent.....	77
5.13.2 Discarding changes to a fabric intent.....	77
5.13.3 Creating a new version of a fabric intent.....	77
5.13.4 Duplicating a fabric intent.....	78
5.14 Deleting a fabric intent.....	78
5.15 Associating planned nodes with real hardware.....	79
5.16 Deploying a fabric intent.....	80
5.16.1 Adding a fabric intent to the deployment pipeline.....	80
5.16.2 Deploying a fabric intent from the deployment pipeline.....	80
5.16.3 The Deployment Trigger Percentage setting.....	81
5.16.4 Deploying backbone fabric intents.....	82
5.17 Deviations.....	82
5.17.1 Viewing deviations.....	83
5.17.2 Accepting or rejecting deviations.....	84
<b>6 Workload VPN intents.....</b>	<b>86</b>
6.1 Role of a workload intent.....	86
6.2 Elements of a workload intent.....	86
6.2.1 Fabrics.....	87

6.2.2 Subnets.....	87
6.2.3 Sub-interfaces.....	88
6.2.4 QoS profiles.....	88
6.2.5 ACL profiles.....	88
6.3 Viewing a workload intent.....	89
6.3.1 Viewing a workload intent.....	89
6.3.2 Viewing a workload intent as code.....	91
6.3.3 Viewing the event log.....	92
6.4 Creating and managing QoS profiles.....	93
6.4.1 Creating a QoS profile.....	93
6.4.2 Editing a QoS profile.....	94
6.4.3 Deleting a QoS profile.....	94
6.5 Creating and managing match groups.....	95
6.5.1 Creating a match group.....	95
6.5.2 Editing a match group.....	97
6.5.3 Deleting a match group.....	97
6.6 Creating and managing ACL profiles.....	98
6.6.1 Creating an ACL profile.....	98
6.6.2 Editing an ACL profile.....	101
6.6.3 Deleting an ACL profile.....	102
6.7 Creating a workload VPN intent.....	102
6.7.1 Creating the basic workload intent.....	103
6.7.2 Adding subnets to the workload intent.....	104
6.7.3 Adding sub-interfaces to the workload intent.....	105
6.7.4 Configuring BGP.....	107
6.8 Modifying a workload intent.....	109
6.8.1 Editing a workload intent.....	109
6.8.2 Creating a new version of a workload intent.....	110
6.8.3 Updating a group of sub-interfaces by label reference.....	110
6.9 Deleting a workload intent.....	111
6.10 Deploying a workload intent.....	112
6.10.1 Adding a workload intent to the deployment pipeline.....	113
6.10.2 Deploying a workload intent from the deployment pipeline.....	113
<b>7 Maintenance intents.....</b>	<b>115</b>
7.1 Viewing a maintenance intent.....	115

7.1.1 Elements of the maintenance intent Design view.....	115
7.1.2 Viewing affected nodes.....	117
7.1.3 Viewing the event log.....	117
7.2 Creating a maintenance label.....	118
7.3 Labeling objects for maintenance.....	118
7.4 Creating a maintenance intent.....	119
7.5 Duplicating a maintenance intent.....	120
7.6 Deploying a maintenance intent.....	121
7.6.1 Adding a maintenance intent to the deployment pipeline.....	122
7.6.2 Progress of a deployed maintenance intent.....	122
7.6.3 Deploying a maintenance intent from the deployment pipeline.....	123
7.7 Removing a maintenance intent from the deployment pipeline.....	123
7.8 Aborting a maintenance intent.....	123
<b>8 Labels.....</b>	<b>125</b>
8.1 Label types.....	125
8.1.1 Nokia pre-defined labels.....	125
8.1.2 User-configured custom labels.....	126
8.2 The Label Factory.....	127
8.3 Viewing available labels.....	127
8.4 Creating a label.....	128
8.5 User-configured label manipulation.....	128
8.5.1 Editing an existing label.....	128
8.5.2 Deleting an existing label.....	129
8.6 Label assignments to fabric intent elements.....	129
8.6.1 Assigning labels to a fabric intent.....	129
8.6.2 Assigning a label to a specific node in a fabric intent.....	130
8.6.3 Assigning labels to a fabric link.....	131
8.6.4 Assigning labels to an edge link interface.....	131
8.6.5 Removing a label assigned to a fabric intent.....	132
8.6.6 Removing a label assigned to a node in a fabric intent.....	132
8.6.7 Removing a label assigned to a fabric link.....	133
8.6.8 Removing a label assigned to an edge link interface.....	133
8.7 Label assignments to workload intent elements.....	134
8.7.1 Assigning labels to a workload intent.....	134
8.7.2 Assigning a label to a specific sub-interface.....	134

8.7.3 Removing a label assigned to a workload intent.....	135
8.7.4 Removing a label assigned to a specific sub-interface.....	135
8.8 Label assignment management.....	135
8.8.1 Viewing the assignments of a specific label.....	136
8.8.2 Querying the label assignment list.....	136
8.8.3 Creating label assignment list string queries.....	136
<b>9 Inventories.....</b>	<b>138</b>
9.1 Viewing the overall inventory.....	138
9.2 Viewing the inventory of a single fabric intent.....	139
9.3 Viewing more information about a node in the inventory.....	139
9.4 Viewing node states.....	140
9.5 Inventory manipulation.....	141
9.5.1 Editing node information.....	141
9.5.2 Viewing the configuration file for a single node.....	141
9.5.3 Associating individual planned nodes with real-world hardware.....	142
9.5.3.1 Associate a planned node.....	142
9.5.3.2 Associate using the Serial Number column.....	143
9.5.4 Disassociating planned nodes from real hardware.....	143
9.5.5 Changing node associations.....	144
9.5.5.1 Node associations in candidate fabric intents.....	144
9.5.6 Updating the system name of a node in a fabric inventory.....	145
9.5.7 Viewing platform details for nodes in the inventory.....	145
9.5.8 Exporting data about inventory items.....	146
9.5.8.1 Exporting an overall inventory spreadsheet.....	146
9.5.8.2 Exporting association data from the overall inventory.....	147
9.5.8.3 Exporting association data from a fabric intent inventory.....	147
9.5.9 Uploading data about a series of inventory items.....	148
9.5.9.1 Uploading association data to the overall inventory.....	148
9.5.9.2 Uploading association data to the fabric intent inventory.....	148
9.5.9.3 Modifying the system name of nodes in an existing fabric intent.....	148
<b>10 Alarms.....</b>	<b>150</b>
10.1 Alarm list.....	150
10.2 Policies.....	152

<b>11 Operations views</b>	<b>154</b>
11.1 Viewing the operational topology	157
11.2 Viewing operational insights	159
<b>12 Digital Sandbox</b>	<b>160</b>
12.1 Integration with the Fabric Services System	160
12.1.1 Digital Sandbox status display	160
12.2 Creating a region	161
12.2.1 Modifying a region	161
12.3 Fabric intents and the Digital Sandbox	162
12.3.1 Creating a fabric intent	162
12.3.2 Updating the Digital Sandbox	163
12.3.3 Deploying a fabric intent	163
12.4 Workload intents	164
12.4.1 Creating a workload VPN intent for the Digital Sandbox	164
12.4.2 Deploying a workload VPN intent to the Digital Sandbox	164
12.4.3 Updating the Digital Sandbox	164
<b>13 System administration</b>	<b>166</b>
13.1 Application settings	166
13.1.1 Software and image catalogs	166
13.1.1.1 Adding a new software image	167
13.1.2 Global configuration settings	167
13.1.2.1 FTP	167
13.1.2.2 NTP	168
13.1.3 Common application settings	169
13.1.3.1 Geomap	169
13.2 User and resource management	169
13.2.1 Roles	170
13.2.1.1 Viewing a list of existing roles	170
13.2.1.2 Predefined roles	171
13.2.1.3 Creating a role	171
13.2.1.4 Modifying the resource access permissions of a role	172
13.2.1.5 Deleting a role	172
13.2.2 Resource groups	172

13.2.2.1 Viewing a list of resource groups.....	173
13.2.2.2 Predefined resource groups.....	173
13.2.3 User groups.....	174
13.2.3.1 Viewing a list of existing user groups.....	175
13.2.3.2 Predefined user groups.....	175
13.2.3.3 Creating a user group.....	175
13.2.3.4 Assigning the specified role of a user group.....	176
13.2.3.5 Deleting a user group.....	176
13.2.4 Users.....	177
13.2.4.1 Viewing a list of existing users.....	177
13.2.4.2 Creating a new user.....	177
13.2.4.3 Assigning a user to a user group.....	178
13.2.4.4 Assigning a role to a user.....	178
13.2.4.5 Deleting a user.....	178
13.3 Backup and restore.....	179
13.3.1 Backing up.....	181
13.3.2 Restoring a backup and installing the Fabric Services System application.....	182
13.3.2.1 Setting up new Fabric Services System nodes.....	183
13.3.2.2 Installing the Kubernetes cluster.....	183
13.3.2.3 Restoring a backup.....	184
13.3.2.4 Installing the Fabric Services System application.....	185
13.4 Capturing troubleshooting data.....	186
13.4.1 How to run the technical support script.....	186
13.5 Node discovery.....	187
13.5.1 Participants.....	187
13.5.2 The discovery and configuration process.....	188
13.5.3 Using an external DHCP server.....	189
<b>Appendix: Fabric intent parameters.....</b>	<b>190</b>
Required parameters.....	190
Optional parameters.....	193
<b>Appendix: Workload VPN intent parameters.....</b>	<b>195</b>
Workload intent.....	195
<b>Appendix: Supported alarms.....</b>	<b>198</b>

# 1 About this document

This Fabric Services System *User Guide* describes the system's user interface (UI), and includes procedures that guide you through the design and deployment of a fabric intent.

This document is intended for network technicians, administrators, operators, service providers, and others who use the Fabric Services System.

**Note:**

This manual covers the current release and may also contain some content that will be released in later maintenance loads. See the *Fabric Services System Release Notes* for information about features supported in each load.

## 1.1 What's new

This is the first release of this document. In future releases, a table will define new or change information for the release.

## 1.2 Precautionary and information messages

The following are information symbols used in the documentation.



**DANGER:** Danger warns that the described activity or situation may result in serious personal injury or death. An electric shock hazard could exist. Before you begin work on this equipment, be aware of hazards involving electrical circuitry, be familiar with networking environments, and implement accident prevention procedures.



**WARNING:** Warning indicates that the described activity or situation may, or will, cause equipment damage, serious performance problems, or loss of data.



**Caution:** Caution indicates that the described activity or situation may reduce your component or system performance.



**Note:** Note provides additional operational information.



**Tip:** Tip provides suggestions for use or best practices.

## 1.3 Conventions

Commands use the following conventions

- **Bold** type indicates a command that the user must enter.

- 
- Input and output examples are displayed in Courier text.
  - An open right angle bracket indicates a progression of menu choices or simple command sequence (often selected from a user interface). Example: **start > connect to**
  - Angle brackets (< >) indicate an item that is not used verbatim. For example, for the command **show ethernet <name>**, **name** should be replaced with the name of the interface.
  - A vertical bar (|) indicates a mutually exclusive argument.
  - Square brackets ([ ]) indicate optional elements.
  - Braces ({ }) indicate a required choice. When braces are contained within square brackets, they indicate a required choice within an optional element.
  - *Italic* type indicates a variable.

Examples use generic IP addresses. Replace these with the appropriate IP addresses used in your system.

## 2 Overview

The Fabric Services System supports intent-based automation for the management of data center fabrics, including:

- planning, designing, and validating a prospective fabric. These are sometimes called "Day 0" activities, and precede actual deployment.
- deploying a planned fabric configuration to hardware for the first time; also, designating resources within your deployed fabric for allocation to specific sources of demand upon its traffic and processing capacity. These are sometimes called "Day 1" activities, and pertain to the deployment itself.
- telemetry monitoring, state monitoring, and subsequent re-configuration of a fabric. These are sometimes called "Day 2+" activities, and are performed some time after the initial deployment.

This *User Guide* describes how to use the system's Graphical User Interface (GUI) to design, test, and manage your data center fabric and the to manage the workloads that place storage and processing demands upon them.

### 2.1 Concepts

The Fabric Services System uses the following concepts for describing and managing the elements of a fabric:

- **Fabric:** a group of switches that are managed as a single logical unit. A single data center can include many complementary and mutually supporting fabrics. These fabrics form an structured "underlay" onto which connections and services can be superimposed.

Fabrics are designed in a hierarchical structure. Typically this structure extends:

- from one or more backbone nodes, providing a gateway that manages communication externally to a WAN
- to one or more spine nodes that aggregate and distribute traffic
- to a collection of leaf nodes each associated with one or more spine nodes
- and from those leaf nodes through edge link interfaces to termination points

Although that description is typical, the presence or absence, size, and role of each level varies widely in actual fabric designs.

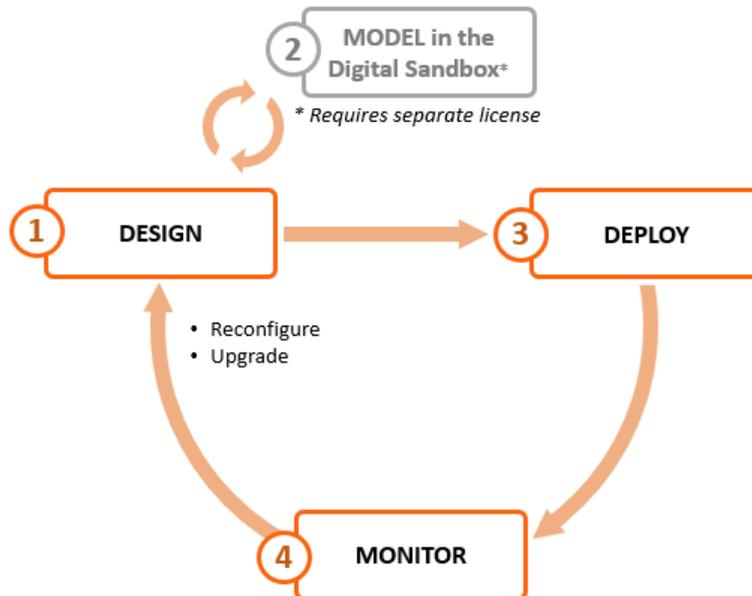
- **Fabric intent:** a set of node configurations that, when deployed to hardware, result in a functioning fabric. An intent gathers and deploys all of the required node configuration tasks into a single transaction, which then deploy successfully or not at all. If the deployment of any participating node's configuration fails at any point, the entire transaction is rolled back to restore all of the participating nodes to their state before the deployment of the intent began.
- **Workload:** a single source of demand upon a data center's fabrics. For example, all of the traffic from a single customer, or tenant, could be directed to only those resources encompassed by a particular workload intent. A workload is an overlay that can be superimposed upon the fabric's previously configured underlay structure.

- **Workload intent:** a set of node configurations pertaining to a specific subset of fabric resources (one or more fabrics, subnets, and sub-interfaces). After the workload intent is defined and its configuration data has been deployed to participating nodes, the fabric resources encompassed by a workload intent can be made available to manage a particular workload. Like a fabric intent, The Fabric Services System deploys a workload intent to all participating nodes as a single transaction. This ensures that the entire deployment succeeds completely, or else fails completely and cleanly, leaving nodes in their pre-deployment states.

## 2.2 Life cycle of a fabric

The Fabric Services System provides tools to assist you in designing, deploying, and managing your data center fabric using the high-level cycle shown in [Figure 1: Fabric life cycle](#).

Figure 1: Fabric life cycle



### 2.2.1 Design the fabric intent

Using the Fabric Services System, you can design a new fabric with a few clicks. The system then creates a detailed fabric intent that identifies all of the necessary nodes and includes all of the configuration files necessary to provision those nodes when they are available.

You can review and modify this fabric intent before proceeding with deployment, using the system's built-in deployment system to reconfigure participating nodes.

During design, you indicate whether you are designing a fabric intent for deployment to real hardware, or to the Fabric Services System Digital Sandbox, a simulator where you can validate prospective fabric designs. The Digital Sandbox is described in [Model in the Digital Sandbox](#).

Figure 2: Design workflow



To create the initial design for your fabric intent, you first select a Nokia certified fabric template and then configure some basic parameters. The system then generates a recommended topology for the fabric you described. The procedure for creating a fabric intent is described in [Fabric intents](#).

After creating your fabric intent, you can (optionally) explore the fabric topology and view details about its constituent nodes and links. The procedure for viewing these details is described in [Fabric intents](#).

Labels are tags that help you group and organize fabric objects according to specific criteria. You can create labels in the Label Factory and then assign them to fabric intents or to objects within a fabric intent.

Labels are a powerful tool you can use to identify a group of objects that can then be subject to collective actions, such as software updates. For example, you can create a label and apply it to a group of items that are subject to a particular upgrade. Then, you can apply the upgrade to all of the nodes with that label as a single action, instead of upgrading them individually.

For the procedures to create and assign labels, see [Labels](#).

## 2.2.2 Model in the Digital Sandbox

If you have purchased the license for the Fabric Services System's Digital Sandbox, you can deploy your fabric to a virtual environment to evaluate its design without relying on physical hardware.

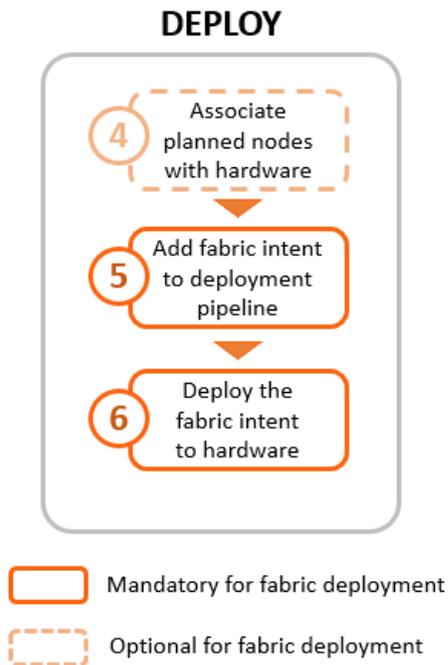
The Digital Sandbox is a network simulator that can emulate data center fabric designs (underlays) and the workload constraints configured upon those fabrics (overlays). Each SR Linux node within the fabric is emulated as its own virtual machine within the cluster, running its own copy of the SR Linux operating system.

In its current form, the Digital Sandbox can emulate a region, the structures of fabrics within that region, and the workload constraints that are configured upon those fabrics (including the edge links that are referred to by the workload). It does not yet simulate dynamic features like traffic flow between the simulated nodes and their endpoints.

### 2.2.3 Deploy the fabric intent

When you are satisfied with the fabric design, you can use the system to deploy your fabric intent to physical hardware.

Figure 3: Deploy workflow



Before deployment, you can associate any planned node in the fabric intent with its available, physical counterpart by providing the corresponding hardware serial number. You can associate individual nodes directly in the UI, or associate nodes in bulk by exporting a spreadsheet from the fabric intent, adding the necessary information for each node, and then re-importing the spreadsheet.

For the procedure to associate nodes individually or in bulk, see [Inventories](#).

Whether you have performed this association or not, you can then add the fabric intent to the deployment pipeline. This pipeline manages all of the intents awaiting deployment, of which there can be many when multiple operators are simultaneously designing and maintaining multiple fabrics.

After the fabric intent is in the deployment pipeline, you can deploy the fabric. The system deploys configuration data to all of the participating, available nodes as a single transaction. The procedure for deployment is described in [Fabric intents](#).

Even if none of the fabric's planned nodes are available when you design your fabric intent, you can still proceed with deployment. The system can deploy the planned fabric onto nodes as they become available; the fabric intent persists in the deployment pipeline until the entire deployment is complete. This allows you to plan a large fabric at the outset, and then deploy it in segments as the hardware become available over time.

## 2.2.4 Monitor the fabric intent

You can use the Operations features of the Fabric Services System to monitor the behavior of your fabric. These tools clearly signal any failure or design deviation, enabling you to take remedial steps as quickly as possible.

## 2.2.5 Repeat the cycle

As node replacement or software updates become necessary, you can use the system's design capabilities to again plan the necessary changes as a new version of the existing fabric intent. You can also model the planned update in the Digital Sandbox before deployment.

When the updated design is satisfactory, you can use the system's deployment capabilities to commit the changes to your working fabric.

You can repeat this cycle as often as needed throughout the lifetime of your fabric.

## 2.3 Workload VPN intents

The fabrics you deploy within a data center can carry traffic belonging to many customers (or "tenants"), all of which impose their own traffic, processing, and storage demands.

The Fabric Services System allows you to define a set of fabric resources that can be used to support an individual workload. This distributes the traffic and processing load efficiently over your fabric, and ensures that a tenant has reliable access to their share of the fabric's capacity. The tool that the Fabric Services System uses to allocate fabric resources is the workload intent.

A workload intent can encompass one or more fabrics. It identifies a set of subnets and their sub-interfaces within the participating fabrics to be made available to a particular source of demand. Sub-interfaces can be configured with Quality of Service (QoS) settings to prioritize traffic and Access Control Lists (ACLs) to accept or reject traffic originating from, or headed to, specific IP addresses.

The Fabric Services System uses the concept of an intent to encompass the set of node configurations that embody the workload intent. The system deploys the workload intent to all participating nodes as a single transaction. This approach ensures that the entire deployment succeeds completely, or else fails completely and cleanly to facilitate another attempt.

Procedures for creating and deploying a workload VPN intent can be found in [Workload VPN intents](#).

## 2.4 Maintenance intents

A maintenance intent identifies a set of nodes and a configuration change that the Fabric Services System deploys to those nodes.

The Fabric Services System supports two types of maintenance intent:

- software changes, which upgrades or downgrades the SR Linux software version running on one or more nodes within a single fabric.
- node replacement, which updates the hardware association for an existing node to a new piece of matching hardware with a different serial number, and downloads the necessary configuration files to the new node so that it can fully resume the role of the hardware it replaced.

Procedures for creating and deploying a maintenance intent can be found in [Maintenance intents](#).

## 2.5 Additional resources

You can obtain more information about the Fabric Services System user interface and its API from the other documents in the Fabric Services System documentation suite.

For the API, more information is available from the Swagger interface directly in the Fabric Services System UI.

### 2.5.1 Documentation resources

Table 1: Fabric Services System information sources

Document	Description
<i>Fabric Services System Release Notes</i>	Release notes accompany every release of the Fabric Services System. These highlight the most up-to-date information about supported features, supported hardware, known limitations, and resolved issues.
<i>Fabric Services System Software Installation Guide</i>	The <i>Software Installation Guide</i> documents the system requirements for the Fabric Services System and detailed procedures for obtaining and installing the Fabric Services System software.
Fabric Services System contextual help	Contextual help is available for many pages within the Fabric Services System UI. To view contextual help for the current page, click  near the upper right of the page.
<i>Fabric Services System User Guide</i>	The <i>User Guide</i> is the primary source for conceptual and practical information about everything you can do in the Fabric Services System GUI. It explores the concepts behind the Fabric Services System, describes the common features available throughout the UI, and provides detailed procedures you can follow to perform all of the main operations available within the system.

### 2.5.2 Accessing the API descriptions from the system GUI

#### About this task

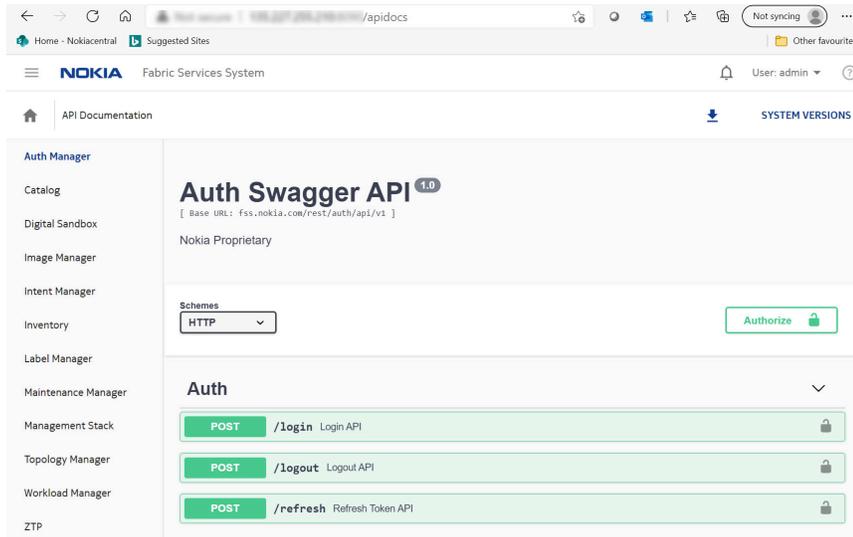
You can view information about API calls directly in the Fabric Services System GUI at any time. This view provides you with the most up-to-date list of APIs for your system and release.

From the system GUI, you can view currently supported APIs, read a description of their purpose, and test API calls using the Swagger interface.

#### Procedure

**Step 1.** Sign in to the Fabric Services System.

**Step 2.** In your browser's URL field, enter the following to open the Swagger API interface: *<GUI URL>/apidocs*



## 3 User interface basics

The Fabric Services System graphical user interface (GUI) allows you to create, view, deploy, monitor, and update fabric intents, workload intents, and maintenance intents. It also allows you to configure supporting systems that enable the management of those intents.

This section describes common operations in the user interface, and how to use some elements that recur throughout the GUI.

### 3.1 Signing in to the Fabric Services System

#### About this task

Follow this procedure to sign into the Fabric Services System user interface using your browser.

#### Procedure

- Step 1.** Open the login page by navigating to the following address in your browser: `http://fsp.nokia.com:8090/login`
- Step 2.** Enter a Username and Password.
- Step 3.** Optionally, select the **Remember username** check reuse the same credentials for your next sign-in.
- Step 4.** Click **SIGN IN**.

The first time you sign in to the Fabric Services System, a welcome dialog displays. This dialog advises you to begin by creating a region to contain fabric intents and workload intents. This dialog does not display in subsequent sessions, even if you have not yet created a region.

#### Expected outcome

The Dashboard displays.

### 3.2 Signing out of the Fabric Services System

#### Procedure

- Step 1.** Click the username displayed at the upper right of any page.
- Step 2.** In the resulting drop-down menu, click **SIGN OUT**.

### 3.3 The dashboard

The Fabric Services System dashboard shows high-level information about the fabric intents, workload intents, and maintenance intents in the system, as well as information about the deployment pipeline.

Here you can see important status information about the intents and deployment pipeline at a glance, and the display can immediately guide you to critical issues that need your attention.

From the dashboard, you can navigate to anywhere within the system.

Figure 4: Fabric Services System dashboard

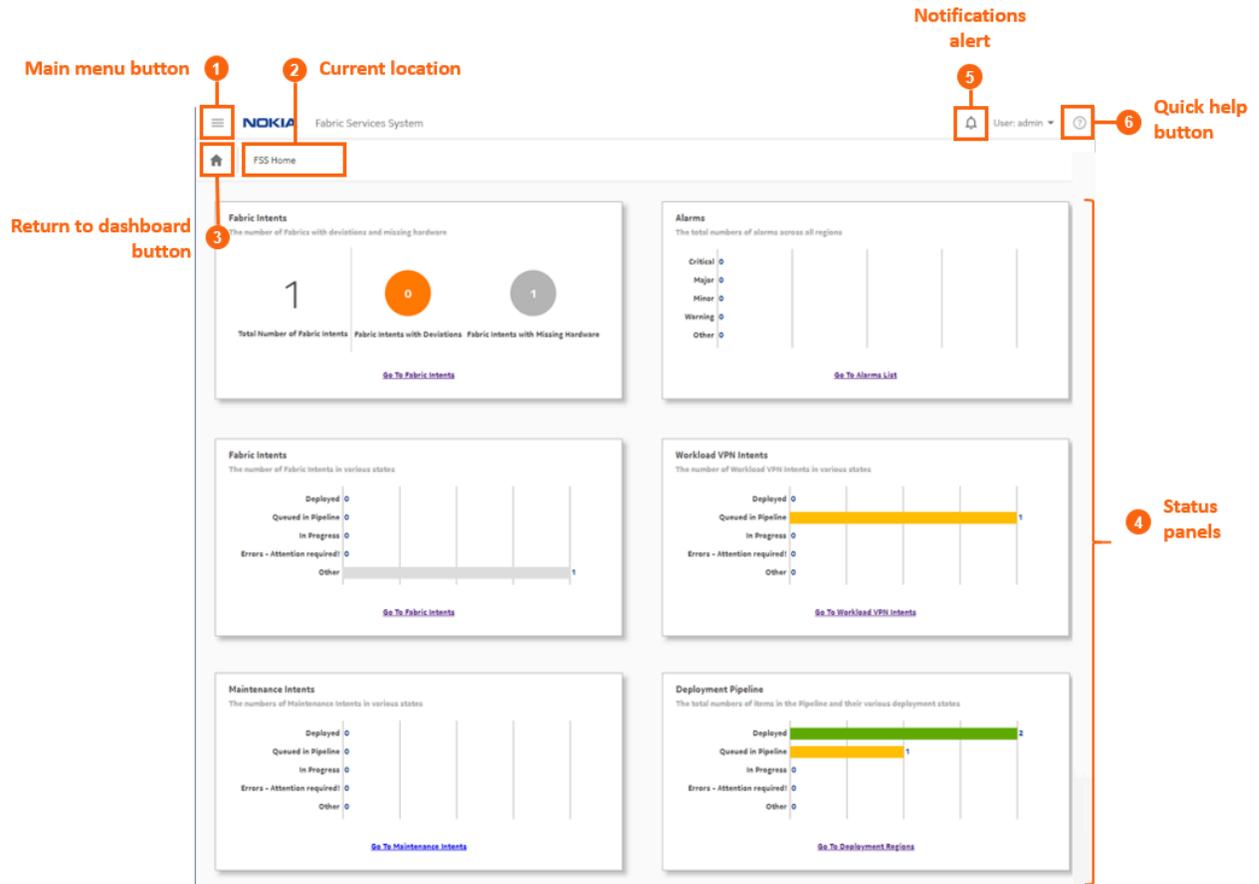


Figure 5: Dashboard deployment regions map

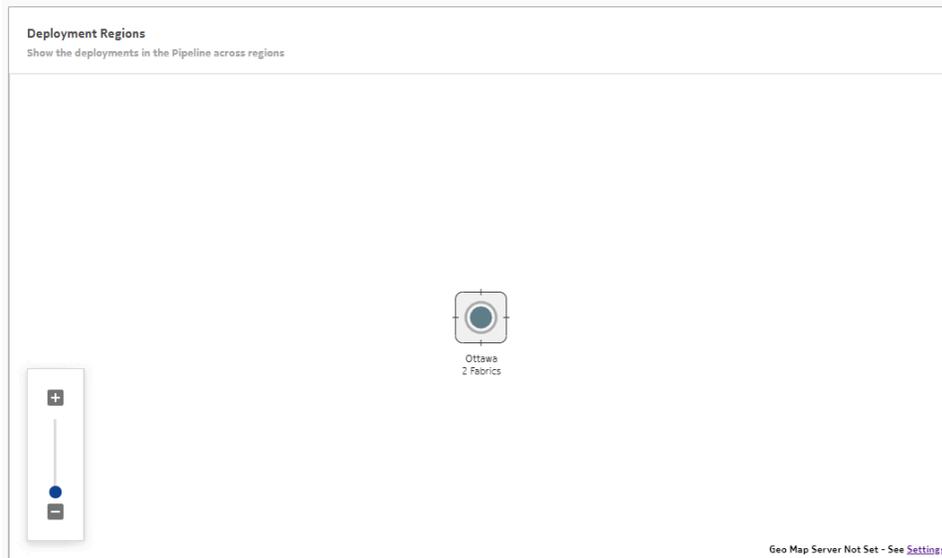


Table 2: Elements of the dashboard display

#	Description
1	The <b>Main menu button</b> opens the navigation panel on the left side of the display. See <a href="#">The main menu</a> for more information about this menu.
2	The <b>Current location</b> indicator displays the name of the current page.
3	The <b>Return to dashboard</b> button takes you directly back to the dashboard. This button displays on all screens within the UI.
4	<p>Status panels</p> <p>The first <b>Fabric Intents</b> panel displays a summary of the fabrics that are:</p> <ul style="list-style-type: none"> <li>the subject of deviations (the configuration of at least one participating node has been modified by an external user since the fabric was deployed)</li> <li>missing hardware</li> </ul> <p>A link at the bottom of the panel opens the Fabric Intents page, which shows a list of all fabric intents from which you can view full details about each intent in the list.</p> <p>Fabric intents are fully described in <a href="#">Fabric intents</a>.</p> <p>The <b>Alarms</b> panel indicates the number of current alarms of each possible severity level (Critical, Major, Minor, Warning, Other).</p> <p>A link at the bottom of the panel opens the Alarms List page, which shows a list of all alarms and from which you can view full details about each alarm in the list.</p>

#	Description
	Fabric intents are fully described in <a href="#">Alarms</a> .
	<p>The second <b>Fabric Intents</b> panel displays a chart showing the number of fabric intents in each of the following states.</p> <ul style="list-style-type: none"> <li>• <b>Deployed:</b> the fabric intent configuration has been deployed to participating nodes, and is in a functioning state with no configuration changes.</li> <li>• <b>In Queue:</b> the fabric intent has been added to the deployment queue, but is not yet deployed. It may be waiting for a user action to begin deployment, or may be waiting until preceding deployments are complete. The system does not begin a new deployment until the preceding intents have been fully deployed.</li> <li>• <b>In Progress:</b> the system is deploying the fabric intent but it is not complete.</li> <li>• <b>Error:</b> there is an error state that requires resolution. Consult the information panel for the affected fabric intent for additional details.</li> <li>• <b>Other:</b> The fabric is in some other state that likely requires attention. Consult the information panel for the affected fabric intent for additional details.</li> </ul> <p>A link at the bottom of the panel opens the Fabric Intents page, which shows a list of all fabric intents from which you can view full details about each intent in the list.</p> <p>Fabric intents are fully described in <a href="#">Fabric intents</a>.</p>
	<p>The <b>Workload VPN Intents</b> panel displays a chart showing the number of workload VPN intents in each of a series of states.</p> <ul style="list-style-type: none"> <li>• <b>Deployed:</b> the workload VPN intent configuration has been deployed to participating nodes, and is in a functioning state.</li> <li>• <b>In Queue:</b> the workload VPN intent has been added to the deployment queue, but is not yet deployed. It may be waiting for a user action to begin deployment, or may be waiting until preceding deployments are complete. The system does not begin a new deployment until the preceding intents have been fully deployed.</li> <li>• <b>In Progress:</b> the system is deploying the workload VPN intent but it is not complete.</li> <li>• <b>Error:</b> there is an error state that requires resolution. Consult the information panel for the affected workload VPN intent for additional details.</li> <li>• <b>Other:</b> the workload is in some other state that likely requires attention. Consult the information panel for the affected workload VPN intent for additional details.</li> </ul> <p>A link at the bottom of the panel opens the Workload VPN Intents page, which shows a list of all workload VPN intents and from which you can view full details about each intent in the list.</p> <p>Fabric intents are fully described in <a href="#">Workload VPN intents</a>.</p>
	<p>The <b>Maintenance Intents</b> panel displays a chart showing the number of maintenance intents in each of the following states:</p>

#	Description
	<ul style="list-style-type: none"> <li>• <b>Complete:</b> the maintenance intent configuration has been deployed to participating nodes, and is in a functioning state.</li> <li>• <b>In Queue:</b> the maintenance intent has been added to the deployment queue, but is not yet deployed. It may be waiting for a user action to begin deployment, or may be waiting until preceding deployments are complete. The system does not begin a new deployment until the preceding intents have been fully deployed.</li> <li>• <b>In Progress:</b> the system is deploying the maintenance intent but it is not complete.</li> <li>• <b>Error:</b> there is an error state that requires resolution. Consult the information panel for the affected maintenance intent for additional details.</li> <li>• <b>Other:</b> in some other state that likely requires attention. Consult the information panel for the affected maintenance intent for additional details.</li> </ul> <p>A link at the bottom of the panel opens the Maintenance Intents page, which shows a list of all maintenance intents and from which you can view full details about each intent in the list.</p> <p>Maintenance intents are fully described in <a href="#">Maintenance intents</a>.</p>
	<p>The <b>Deployment Pipeline</b> panel displays a chart showing the number of intents that are currently in the system's deployment pipeline and are in each of the states described previously for fabric, workload, and maintenance intents.</p> <p>A link at the bottom of the panel opens the Deployment Regions page, from which you can view the complete list of intents in the pipeline with additional details.</p> <p>Deployment regions, and the deployment pipeline, are fully described in <a href="#">Deployment regions</a>.</p>
5	<p>The <b>Notification alert</b> indicates whether there are active notifications pertaining to the Fabric Services System. When it is active, click this icon to see the active notifications.</p>
6	<p>The <b>Quick Help</b> (  ) button displays help topics pertaining to the current screen.</p>
7	<p><b>Deployment regions map:</b> below the information panels, the system displays a geographic map showing the existing deployment regions. Currently, the Fabric Services System supports the creation of only one deployment region.</p> <p>Double-click a region icon to open the Regions page. From there, you can view more information about the selected region.</p> <p>This map works identically to the region map described in <a href="#">Deployment regions</a>.</p>

### 3.3.1 Notifications

Notifications are alerts generated by different subsystems within the Fabric Services System. When a new notification is raised, the bell icon at the upper right of the UI displays a red circle. A counter beside the icon indicates how many new notifications have accumulated since you last viewed the list of individual items.

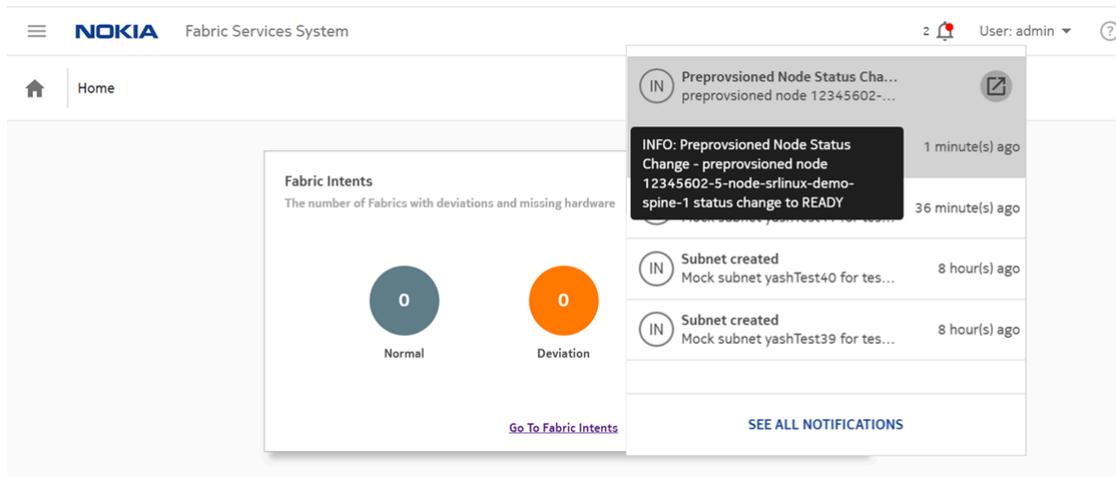
Figure 6: Notifications indicator



Clicking the bell icon displays the list of individual notifications. Hovering over a notification in the list displays additional details about that item.

Clicking on the notification takes you to the page in the UI that corresponds to the subsystem that triggered the notification. For example, in the illustration below, the inventory subsystem has triggered a notification when the status of a node changed from "READY" to "NOT\_READY". The link in the notification takes you to the global inventory page.

Figure 7: Notifications list



Within the notifications list, items are sorted by order of occurrence with the most recent notifications displayed at the top of the list.

#### Notification types

The system currently generates a notification for any device that changes its status to Not Ready.

#### Notification severity

The system currently supports a single severity level: Alert.

## 3.4 The main menu

The main menu button (☰) opens a set of links to the main pages of the Fabric Services System. You can use the main menu to access major features from anywhere within the UI.

The main menu is divided into the following categories:

- **INTENT**, which contains links to pages where you can create and manage Region, Fabric, Workload, and Maintenance intents.
- **OPERATIONS**, which contains links to monitor and maintain fabrics.
- **COMMON RESOURCES**, which contains links to supporting systems like profiles, the Label Factory, and the hardware inventory.
- **Settings**, for managing the system, authentication, and communications.

### INTENT links

Intent links consist of:

- **Deployment Regions:** this page shows a logical view of the fabrics managed by the system, divided first into regions, and then into smaller organizational units. From this page you can:
  - create, delete, or edit a region
  - open a region to view its topology and other details
  - access the deployment pipeline, from which you can deploy intents previously added to the pipeline
- **Fabric Intents:** this page shows a list of all fabric intents regardless of state, and includes basic information about each intent. From this page you can:
  - view a list of fabric intents
  - begin creating a new fabric intent
  - open a fabric intent to view its topology and other details
  - duplicate or create a new version of an existing fabric intent
  - manage the edge links associated with a fabric intent, including creating a LAG or configuring a breakout port
  - delete a fabric intent
  - add a fabric intent to the deployment pipeline
- **Workload VPN Intents:** this page shows a list of all workload VPN intents, and includes basic information about each intent. From this page, you can:
  - view a list of workload intents
  - begin creating a new workload intent
  - open a workload intent in its own page to view additional details
  - delete an existing workload intent
  - add a workload intent to the deployment pipeline
- **Maintenance Intents:** this page shows a list of all maintenance intents, and includes basic information about each intent. From this page, you can:
  - begin creating a new maintenance intent
  - duplicate an existing maintenance intent

- open a maintenance intent in its own page to view additional details
- delete an existing maintenance intent
- add a maintenance intent to the deployment pipeline
- abort a maintenance intent that is underway

## OPERATIONS links

Operations links consist of:

- **Operational and Health Insights:** this page displays a map of configured regions, and allows you to drill down to the fabric and node level to view health status and any deviations for displayed objects.
- **Alarms List:** this page displays information about all alarms currently affecting managed fabrics and the objects within them.
- **Logs:** opens a separate Kibana tab to show system microservice data, if a Kibana URL has been configured. See [Application settings](#).

## COMMON RESOURCES links

- **Inventory:** this page shows a list of all hardware that the system is aware of. From this page you can:
  - add or delete hardware
  - associate planned nodes with real hardware when available
- **Label Factory:** this page shows a list of all labels that are available for assignment to objects displayed in the UI. From this page, you can:
  - create or delete labels
  - see which labels have already been assigned, and to which objects
- **Profiles:** this page displays a list of profiles available to the system for each of the following (on separate tabs):
  - Quality of Service (QoS)
  - Access Control List (ACL)
  - Match Groups
- **Policies:** this page allows you to customize the severity level for individual alarms.
- **User and Resource Management:** this page displays the Fabric Services System user, group, and permissions data.

## Settings

The Settings link opens a page from which you can configure some aspects of the system's behavior with regard to:

- **Image Management:** an interface for adding the SR Linux images to the system library
- **Software Catalog:** a list of SR Linux images currently in the system library
- **Global Configuration:** an interface for managing FTP and NTP settings
- **Common Application Settings:** an interface for managing the image displayed as a map background, and external services.

### 3.5 Detail panels

Many pages in the Fabric Services System UI allow you to open panels on the right side of the page to display more information. This could be information about the overall display, or about selected objects.

Table 3: Detail panels

Icon	Panel	Description
	Deviations	For a fabric intent, this panel shows a cumulative list of deviations logged for the fabric intent. Each item in the deviations list identifies a configuration change originating from outside the Fabric Services system. For more information about deviations, see <a href="#">Deviations</a> .
	Information	This panel displays details about the object that is currently selected in the main display area.
	Legend	This panel displays the different shading used to highlight elements within the fabric intent, and their significance. The same panel also displays the set of icons that can appear in the main display area, and their significance.
	Errors	For a fabric intent, this panel displays a list of errors that require resolution. Double-click an error in the list for more information about that error that can guide troubleshooting. Controls in the Errors panel allow you to filter and sort the displayed list.
	Needs Attention	For a fabric intent, this panel displays a list of status issues that are not errors or deviations, but may require resolution. Double-click an item for more information about that item that can guide troubleshooting. Controls in the Needs Attention panel allow you to filter and sort the displayed list.

### 3.6 Lists

Many pages in the Fabric Services System UI display lists of objects. Wherever they appear, these lists include standard controls to help you sort and filter lists to focus on the information you need.

Figure 8: List controls

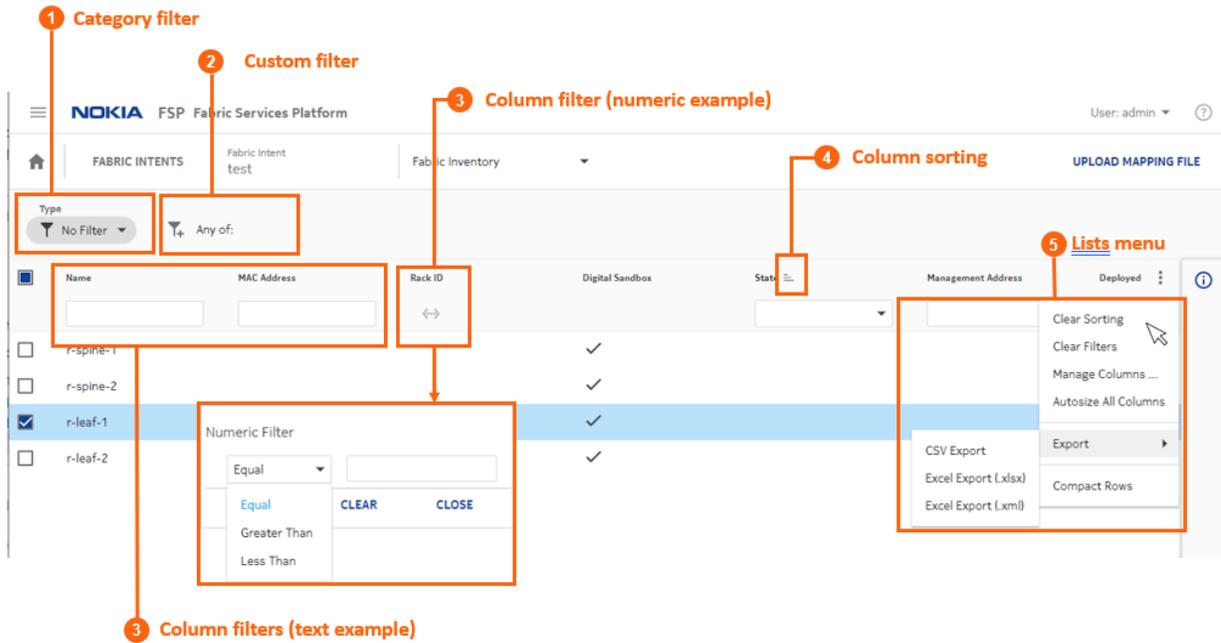


Table 4: List controls

#	Control type	Description
1	Category filter	Some pages in UI can display two entirely different types of the same data. For example, the Inventory page can display a list of either physical hardware, or virtual hardware used in fabric simulations.  The category field allows you to indicate which type of data you want the page to display.
2	Custom filters	You can build custom filters with multiple parameters using this control, as described in <a href="#">Custom filters</a> .
3	Column filters	You can constrain the list to show only items that match the values you specify in one or more column filters.  Filtering is not supported for columns of check box values. <ul style="list-style-type: none"> <li>• If the column filter is a field, type letters and numbers that you want matched in the values in that column.</li> <li>• If the column filter is a drop-down list, select a value that you want matched in the values in that column.</li> <li>• If the column filter is a calendar control, click the calendar icon and select a date that you want matched in the date values in that column.</li> </ul>

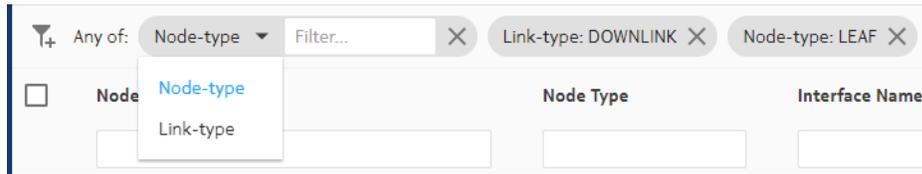
#	Control type	Description
		<ul style="list-style-type: none"> <li>If the column filter is a range control, click the range icon and enter minimum and maximum values to define a value range. Values for the column must fall within that range.</li> </ul>
4	Column sorting	<p>Click a column title to cycle between three sorting states:</p> <ul style="list-style-type: none"> <li>Ascending order (  )</li> <li>Descending order (  )</li> <li>No sorting</li> </ul> <p>The system does not support sorting for columns of check box values.</p>
5	Lists Menu	Click the  icon to open a list of common list controls that allows you to control the behavior of most lists in the UI:
	Clear Sorting	Clears all sorting that has been set for all columns of the list.
	Clear Filters	Clears all filters that have been set for all columns of the list.
	Manage Columns...	Opens a list of all columns in the list, which you can then enable (to show) or disable (to hide). These selections persist for subsequent visits to the same page.
	Autosize All Columns	Sets the width of each column in the list to the minimum width required to display the heading without wrapping its text, and at least as wide as the filter text field for that column.
	Export to CSV, Export to XLSX, Export to XML	Exports the entire table to either a Comma-separated Values (CSV) list, a spreadsheet in XLSX format, or a spreadsheet in XML format.
	Compact Rows	Reduces the white space allocated to each row, fitting more rows onto each page.

### 3.6.1 Custom filters

#### About this task

You can define one or more filters to limit the set of rows displayed in any list. The fields on which you can filter, and the values to match in those fields, are determined by the system.

Figure 9: Adding custom filters to a list



To add a custom filter to a list:

### Procedure

- Step 1.** In the panel across the top of the list, Click the **Add Filter** icon (  ).
- Step 2.** Click the filter drop-down and select a field name from the list.
- Step 3.** Click in the **Filter...** field and select a value for the field from the resulting drop-down list.  
This is the value that must be matched in that field in order for a row to qualify for inclusion in the list.
- Step 4.** To add another filter, repeat steps 1 through 3.
- Step 5.** Continue adding filters until your custom filter is as precise as you need.

## 4 Deployment regions

In the Fabric Services System, a deployment region (or just "region") is a container for other intents. A region serves the following purposes:

- it logically groups intents together on the system map to make them more manageable.
- it is the level at which some connection properties are set; these properties are shared among all intents within the region.
- it maintains a deployment pipeline for all intents (fabric intents, workload VPN intents, and maintenance intents) within the region.

You can view the deployment pipeline, trigger deployments, and manage the deployment process from the region's deployment pipeline page. The deployment pipeline is a key feature of a region. Because the Fabric Services System can manage complex data center fabrics encompassing many nodes and multiple locations, there is the potential for operators to concurrently design and deploy a large number of intents (whether fabric intents, workload intents, or maintenance intents). The deployment pipeline ensures that resources are allocated to each intent deployment, in sequence, to prevent simultaneous deployments from interfering with each other.

You cannot create an intent without assigning it to a region. For this reason, you must create a region before deploying any fabric intent, workload intent, or maintenance intent.

The system supports the configuration of a single region that contains all intents.

### 4.1 Region properties

Every region maintains a set of properties that are shared by all intents within that region. [Table 5: Region properties](#) lists these properties.

*Table 5: Region properties*

Property	Description
Name	The text string that identifies this fabric in the UI.
Description	The text string that is displayed in the <b>Additional information</b> panel for a region.
Location	The physical location represented by the region.
System IP Pool	<b>Pool name:</b> set as "default" and cannot be changed.
	<b>IP Type:</b> set as "IPv4" and cannot be changed.
	<b>IP Blocks:</b> Contains one or more CIDR blocks representing IP addresses that can be assigned to the management interfaces of devices managed by the system. Enter these blocks using CIDR notation; for example, 1.1.1.1/24.

Property	Description
	<p>If you need more IP addresses for devices in your fabrics than the current CIDR blocks support, you can modify the region to add more CIDR blocks to the System IP pool.</p>
<p>Inter Switch Link IP Pool</p>	<p><b>Pool name:</b> set as "default" and cannot be changed.</p>
	<p><b>IP Type:</b> set as "IPv4" and cannot be changed.</p>
	<p><b>IP Blocks:</b> Contains one or more CIDR blocks representing IP addresses that can be assigned to inter-switch links between devices in real (not Digital Sandbox) fabric intents throughout this region. Enter these blocks using CIDR notation. For example: 1.1.1.1/24.</p> <p> <b>Note:</b> Each link within a fabric intent requires two IP addresses from this block; one for each endpoint. From the set of links that are possible with the IP addresses in the specified CIDR block, two are reserved for use by the system to represent the network IP address and the broadcast address. These are unavailable for inter-switch links. For example, a CIDR block ending in /26 that can support up to 32 links will in fact support up to 30 links within the fabric intents in this region.</p> <p>If you create a fabric intent that requires more links than are available with the current pool, fabric generation fails. The event log for the fabric intent indicates that there are insufficient IP addresses for the required links, and shows the number of addresses required versus the number available.</p> <p>Should you need more IP addresses available to your fabrics than the current CIDR blocks support, you can modify the region to add more CIDR blocks and thereby support additional links.</p>
<p>Out of Band Management IP Pool</p>	<p><b>Pool name:</b> set as "default" and cannot be changed.</p>
	<p><b>IP Type:</b> set as "IPv4" and cannot be changed.</p>
	<p><b>IP Blocks:</b> Contains one or more CIDR blocks representing IP addresses that can be assigned to the primary loopback interface of a device. A router ID can then be derived from the same set of allocated IP addresses. Enter these blocks using CIDR notation; for example: 1.1.1.1/24.</p> <p> <b>Note:</b> From the set of links that are possible with the IP addresses in the specified CIDR block, two are reserved for use by the system to represent the network IP address and the broadcast address. These are unavailable for the out-of-band management IP pool. Should you need more links in your fabrics than the CIDR blocks you specified here support, you can modify the region to add more CIDR blocks to this pool.</p>
<p>BGP ASN Numbers</p>	<p><b>First:</b> the lowest permissible value in a range of Autonomous System Numbers (ASNs) that uniquely identify a network with a unique routing policy.</p>

Property	Description
	The ASN must be unique so that IP address blocks appear to originate from a unique location to which BGP can determine a route. Default value: 400
	<b>Last:</b> the highest permissible value in a range of ASNs, as described above. Default value: 500
Digital Sandbox DHCP Configuration	Digital Sandbox DHCP CIDR: the CIDR block use to assign IP addresses for the SR Linux container nodes within the Fabric Services System Digital Sandbox. For more information about the Digital Sandbox, see <a href="#">Digital Sandbox</a> .
Route Target	<b>Global Index:</b> an underlay service parameter set when workloads are created, which ensures that services created as part of a workload are unique within a region and across multiple regions.
EVPN Profile	<b>Control Plane:</b> the control plane type to use throughout the region. Supported values are: <ul style="list-style-type: none"> <li>• None: no EVPN peering</li> <li>• eBGP</li> <li>• iBGP</li> </ul> Default value: none
Protocol authentication	<b>BGP:</b> Enabled or Disabled Enabling BGP protocol authentication causes the Fabric Services System to create a unique MD5 authentication key. This encrypted key is required to authenticate communication between eBGP peers.  No user action is required to create the authentication key or to incorporate eBGP authentication requiring this key into subsequent node configurations. These actions are handled internally by the system.

## 4.2 Creating a region

### About this task

The Fabric Services System currently supports a single region.

The first time you log into the system, you are prompted to create a region as the first step in creating and managing fabric intents. The prompt does not appear for subsequent logins.

After you have created a region, you cannot create another; however, you can modify the region as described in [Modifying a region](#).

To create a new region:

### Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** In the main menu, select  **Deployment Regions**.

**Step 3.** Click the **+ CREATE A REGION** button.



**Note:** If a region has already been created, this option is disabled.

**Step 4.** Enter values for the parameters described in [Table 5: Region properties](#), up to and including Route Target.

**Step 5.** To configure EVPN peering within the region, define the EVPN profile by doing the following:

- a. If your region will not use EVPN peering, leave the default value of the **Control Plane** drop-down as None and go to step [9](#).
- b. If your region will use eBGP for EVPN peering, select eBGP from the **Control Plane** drop-down list and go to step [9](#).
- c. If your region will use iBGP for EVPN peering, and you plan to use only nodes within a fabric created with the Fabric Services System as route reflectors, you will need to create those fabrics and configure those nodes before you can configure iBGP; go to step [7](#).
- d. If your region will use iBGP for EVPN peering and you plan to use existing nodes (that are not part of a fabric created with the Fabric Services System) as route reflectors, you can configure iBGP; go to step [6](#).

**Step 6.** Configure iBGP by doing the following:

- a. In the Control Plane drop-down, select iBGP.
- b. Enter a valid Autonomous System (AS) value.
- c. Enter a Cluster ID to represent the SR OS cluster to which leaf nodes will be peered.  
The system currently supports a single SR OS cluster for iBGP.
- d. Add a route reflector by clicking **+ADD**, entering a route reflector IP address in the resulting dialog, and clicking **ADD**.
- e. If required for redundancy, add additional route reflector IP addresses using the same method.  
All route reflectors must point to the same SR OS cluster. Leaf nodes can use any of the reflectors in this list interchangeably to communicate with the SR OS cluster.
- f. Go to step [9](#).

**Step 7.** Leave the default value of the **Control Plane** drop-down as None.

After you create a fabric that includes one or more nodes to act as route reflectors, complete [Updating a region to use iBGP](#). As part of that procedure, you change the Control Plane setting to iBGP and select your newly configured route reflectors.

**Step 8.** In the **Protocol Authentication** panel, configure eBGP authentication:

- if eBGP authentication is not required, leave the **BGP** toggle set to Disabled.
- if eBGP authentication is required, set the **BGP** toggle to Enabled. This causes the Fabric Services System to generate a unique key that is required for eBGP authentication and to incorporate the requirement for this authentication into subsequent node configurations.

**Step 9.** Click the **CREATE** button.

## 4.3 Modifying a region

### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** In the main menu, select  **Deployment Regions**. The deployment regions page opens, showing a graphical representation of regions already created.
- Step 3.** Right-click the region and select **Edit Details...** from the contextual menu.
- Step 4.** Modify any of the following settings for the region:
  - General/Description
  - General/Location
  - System IP Pool/IP Blocks
  - Inter Switch Link IP Pool/IP Blocks
  - Out of Band Management Pool/IP Blocks
  - EVPN Profile/Control Plane
  - Protocol Authentication/BGP
- Step 5.** Click **SAVE**. The system saves the new region details.

### 4.3.1 Implications of modifying a region

A region is a container for intents, and every fabric intent you create within a region inherits the region's properties (such as configurations for EVPN peering or eBGP authentication).

If you modify the region's settings, the system does not automatically propagate the changes to pre-existing fabric intents within the region.

To update existing fabric intents with the new region settings, you must regenerate each of the fabric intents.

- If you have saved, but not generated, a fabric intent, no action is required; the region settings are not inserted into the fabric's code until you generate the topology, so the latest settings are picked up immediately when you do generate the fabric topology.
- If you have generated but not deployed a fabric intent, you need to update and re-save the fabric intent, then regenerate the topology. The newly generated fabric intent includes the latest settings from the updated region.
- If you have deployed a fabric intent you need to create a new version of the intent, regenerate the topology, and deploy the result. Regenerating the new version of the fabric intent inserts the new region settings into the fabric code, and deploying the fabric intent updates the working node configurations accordingly.

For the steps to create a new version of an already-deployed fabric intent, see [Creating a new version of a fabric intent](#).

When you regenerate the fabric intent, the system updates the settings to match the current configuration of the region.

You can then add the new version of the fabric intent to the region's deployment pipeline, and deploy the fabric intent from the pipeline normally.

## 4.4 Updating a region to use iBGP

If you plan to use only nodes within a fabric created with the Fabric Services System as route reflectors, you must create that fabric and configure those nodes before you can configure a region to use iBGP for EVPN peering.

When you have created at least one fabric within the region and configured nodes within that fabric as route reflectors, you can return to the Deployment Regions page to update the region settings to use iBGP.

This section assumes you are familiar with the process of creating and deploying fabrics in the Fabric Services System, as described in [Fabric intents](#).

It also assumes you have the expertise to configure nodes as route reflectors, independently of the Fabric Services System.

### 4.4.1 Updating region settings to use iBGP

#### Prerequisites

Before you begin this procedure, you must have created and deployed at least one fabric within the region, and then configured nodes within the region as route reflectors as described in [Configuring fabric nodes as route reflectors](#).

#### About this task

When there are route reflectors in a fabric, you can update the region settings to use those route reflectors for iBGP EVPN peering.

To update the region settings to use iBGP:

#### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** In the main menu, select  **Deployment Regions**. The Deployment Regions page opens, showing a graphical representation of regions already created.
- Step 3.** Select the region icon on the Deployment Regions map, right-click, and select **Edit Details...** from the displayed list.
- Step 4.** In the EVPN Profile section, use the **Control Plane** drop-down to select iBGP.
- Step 5.** Enter a valid Local Autonomous System (AS) value.
- Step 6.** Enter a Cluster ID to represent the SR OS cluster to which leaf nodes are peered. The system currently supports a single SR OS cluster for iBGP.
- Step 7.** Add a previously configured route reflector to the **Route Reflector IP List** by doing the following:
  - a. Click **+ADD**.
  - b. Enter a route reflector IP address in the resulting dialog.
  - c. Click **ADD**.
- Step 8.** If required for redundancy, repeat step 7 to add additional route reflector IP addresses.  
All route reflectors must point to the same SR OS cluster. Leaf nodes use any of the reflectors in this list interchangeably to communicate with the SR OS cluster.

**Step 9.** Click the **SAVE** button. The system saves the new region settings.

### What to do next

You are now ready to proceed to [Updating existing fabric intents to use new iBGP settings](#).

## 4.4.2 Configuring fabric nodes as route reflectors

### Prerequisites

Before you begin this procedure, you must:

- have created a deployment region with EVPN peering set to "None", as described in [Creating a region](#)
- be familiar with the fabric intent creation and deployment process described in [Fabric intents](#)
- know how to configure nodes as route reflectors, independently of the Fabric Services System

### About this task

You can use the Fabric Services System to design and deploy fabrics within a region. You can then configure nodes within those fabrics to function as route reflectors. When there are route reflectors in the fabric, you can update the region settings to use those route reflectors for iBGP EVPN peering.

To configure nodes within a fabric as route reflectors:

### Procedure

- Step 1.** Configure and deploy one or more flexible leaf/spine fabric intents using the Fabric Services System, as described in [Fabric intents](#).
- Step 2.** Note the IP address of each node you plan to configure as a route reflector.
- Step 3.** Manually configure each of these nodes as a route reflector, independently of the Fabric Services System.
- Step 4.** Accept these deviations as described in [Accepting or rejecting deviations](#).

### What to do next

You are now ready to proceed to [Updating region settings to use iBGP](#).

## 4.4.3 Updating existing fabric intents to use new iBGP settings

### About this task

When you have modified the region to use iBGP, you must update all fabrics within the region so that they inherit these new settings. This dependency between a region and its fabrics is described in more detail in [Implications of modifying a region](#).

To update your fabrics with the new iBGP settings, do the following for each fabric in the region:

### Procedure

- Step 1.** Open the fabric intent's Design page.
- Step 2.** Create a new version of the intent by selecting **Create a New Version** from the **More actions**  menu.
- Step 3.** Save the fabric intent.

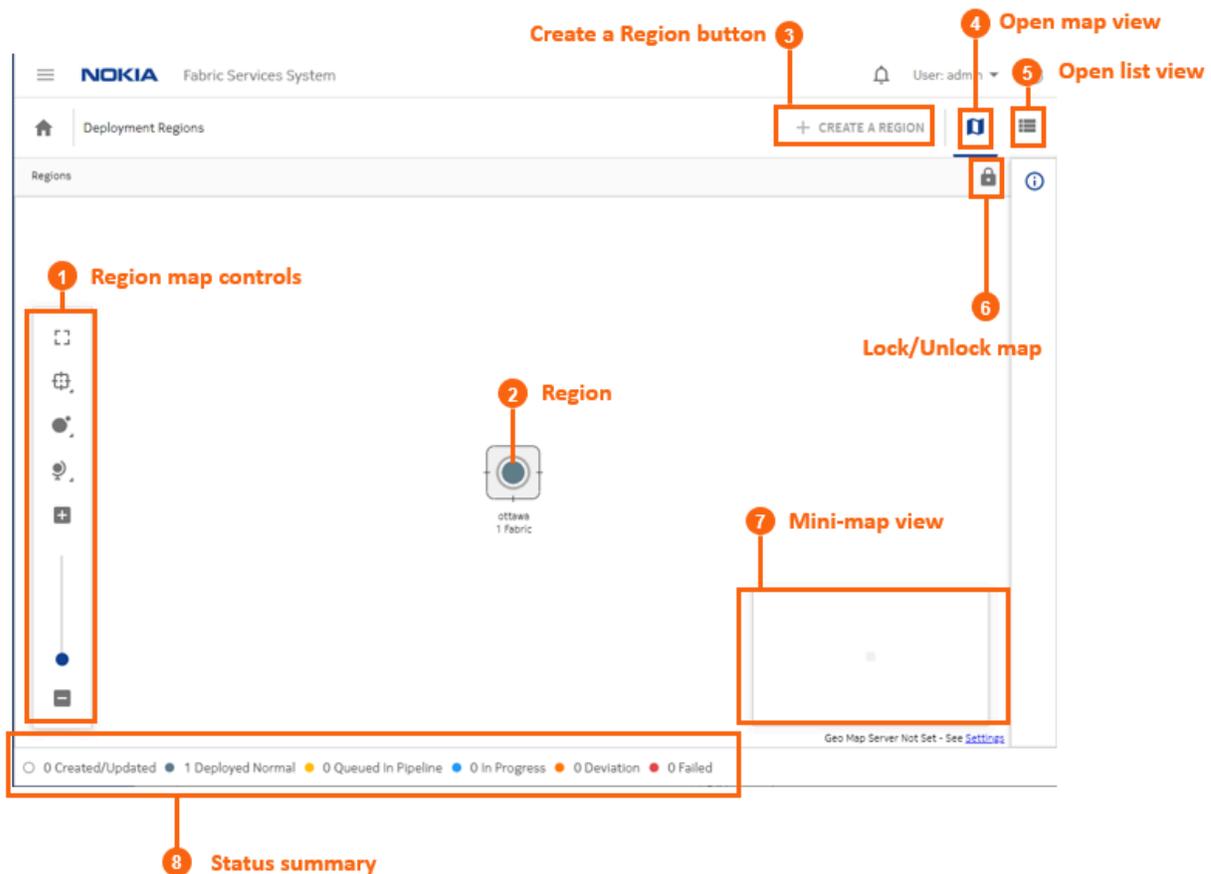
**Step 4.** Click  **GENERATE FABRIC**.

**Step 5.** Deploy the new version of the fabric as described in [Deploying a fabric intent](#).

## 4.5 The region map

The region map is a graphical representation of all regions managed by the system, superimposed over a map image.

Figure 10: The region map



You can change the arrangement of regions on the map, unless the administrator has chosen to lock the map for all users. The region map display is user-specific; your changes to the map persist the next time you log in with the same user account.

An administrator can manage their own version of the map, and periodically deploy it so that the latest version is available to other users.

The background image displayed for the region map is determined by the system's Geomap Tile Server setting. To view or change this setting, go to **Settings** → **Common Application Settings**. For details, see the [Geomap](#) section of [Application settings](#).

After you select an image, it displays the next time you open the region map. You can control the opacity/transparency of the image using the controls described in [Manipulating the region map](#).

## 4.5.1 Viewing and using the region map

### Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** In the main menu, select  **Deployment Regions**. The deployment regions page opens, showing a graphical representation of region already created.

From this map, you can do the following:

- Lock the region icon in its current position by clicking the Lock button (). This replaces the button with an Unlock button ().
- View the set of fabric, workload, and maintenance deployments currently pending for a region by right-clicking on the region and selecting **Show Deployment Pipeline** from the contextual menu.
- View the list of Fabric Intents within a region by right-clicking on the region and selecting **Show Fabric Intent List** from the contextual menu.
- View and optionally edit details about a region by right-clicking on a region and selecting **Edit Details...** from the contextual menu.  
For information about the consequences of such edits, see [Implications of modifying a region](#).
- View information about DHCP settings by selecting **View DHCP Configuration...** from the contextual menu.
- Delete the region by right-clicking on the region and selecting **Delete...** from the contextual menu.  
You cannot delete a region that contains one or more fabrics.

## 4.5.2 Viewing and using the region list

### About this task

Although the region displays as a map by default, you can switch the view to a list of regions instead. However, in the current release, this is a list of one item.

To view a region list instead of a map:

### Procedure

**Step 1.** Open the Regions map as described in [Viewing and using the region map](#).

**Step 2.** Click the **Regional List** () button. The system replaces the map display with a list showing the single region already configured.

From this list you can access most of the options available from the region map by clicking on the **More actions** icon () at the right of the row in the region list and selecting one of the following:

- **Show Deployment Pipeline** to view the set of Fabric, Workload, and Maintenance deployments currently pending for the region.
- **Show Fabric Intent List** to view the list of Fabric Intents within the region.

- **Edit Details...** to view and optionally edit details about a region.  
For information about the consequences of such edits, see [Implications of modifying a region](#).
- **View DHCP Configuration** to view the region's DHCP settings.
- **Delete...** to delete the region.  
You cannot delete a region that contains one or more fabrics.

**Step 3.** To return to the map view from the list view, click the **Regional Map** () button.

### 4.5.3 Manipulating the region map

In addition to being able to drag the region within the map (unless the map has been locked), several controls allow you to manipulate the way elements are displayed on the Regions map.

Table 6: Display controls

Icon	Description
	Lock: prevents you from changing the position of regions on the Regions map. When clicked, the button toggles to Unlock.
	Unlock: restores the ability to move regions on the Regions map. When clicked, the button toggles to Lock.
	Fit to screen: adjusts the magnification setting for the topology display so that the entire topology diagram fits inside the current window.
	Clustering controls: <ul style="list-style-type: none"> <li>• Sets the cluster health display method (no display, or as a pie chart, or as a solid circle).</li> <li>• Enables or disables clustering boundaries.</li> <li>• Enables or disables whether moving a region icon also changes the position of its contents on the underlying map.</li> </ul>
	Vertex controls: adjusts the size of the vertices to either small, medium, large, or extra large.
	Map controls: adjusts the following: <ul style="list-style-type: none"> <li>• Bird's eye view enables or disables the mini-map in the lower right.</li> <li>• Background layer opacity adjusts the background opacity in 5% increments.</li> </ul>
	Zoom in, Zoom out: increases or decreases the magnification for the topology diagram.

## 4.6 The deployment pipeline

The Fabric Services System uses the deployment pipeline to manage the many concurrent fabric, workload, and maintenance intents that can be created when multiple users are working with the system simultaneously.

When you complete the design of a fabric, workload, or maintenance intent and are ready to deploy it, you add the intent to the region's deployment pipeline. The procedures to add intents to the deployment queue are described respectively in [Fabric intents](#), [Workload VPN intents](#), and [Maintenance intents](#).

The intent then remains in the pipeline until you instruct the system to proceed with its deployment.

Figure 11: The Deployment Pipeline

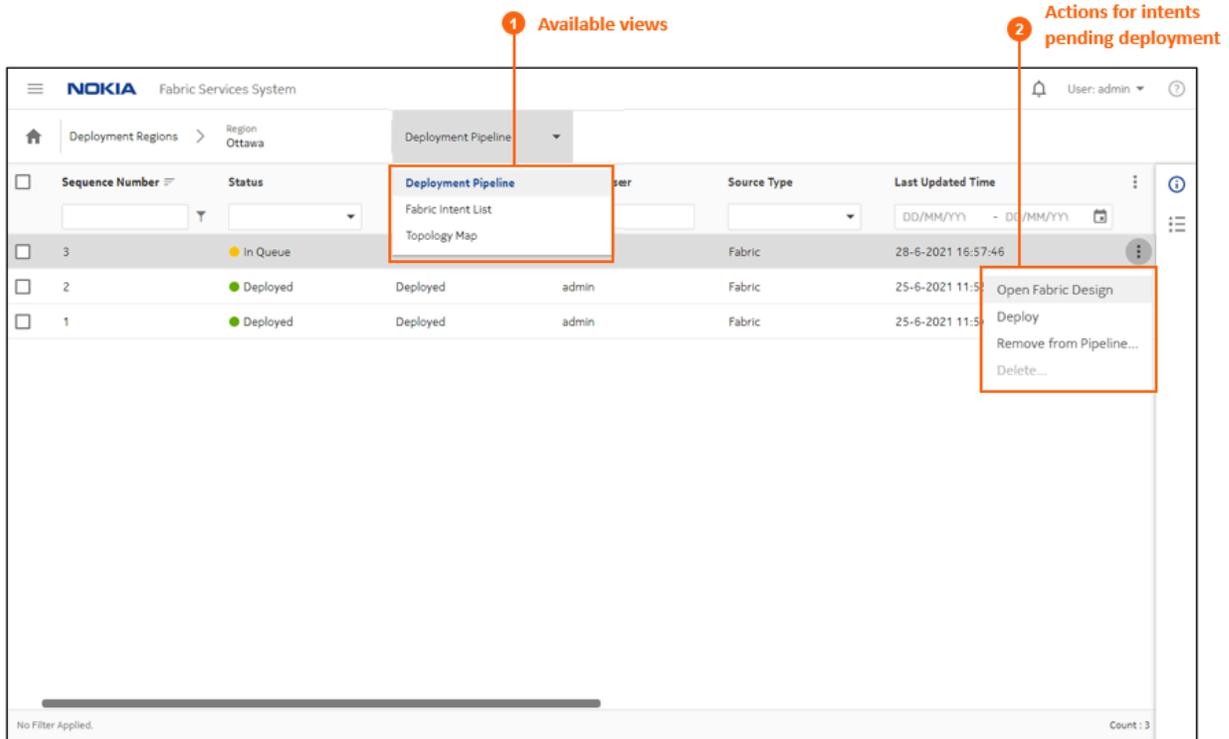


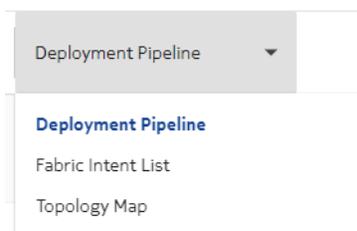
Table 7: Deployment pipeline columns

Column	Description
Sequence Number	The number indicating the sequence (from low to high) in which the displayed intents are deployed within the region. Deployment is currently strictly serial, so one deployment must complete entirely before the next begins.
Status	This indicates the progress the displayed intent has made toward deployment. The system currently supports the following deployment statuses: <ul style="list-style-type: none"> <li>Other: some status not accounted for below.</li> </ul>

Column	Description
	<ul style="list-style-type: none"> <li>• Deployed: deployment of the intent has completed successfully.</li> <li>• Error: deployment of the intent encountered an error and did not complete.</li> <li>• In Queue: the intent is awaiting deployment, likely waiting for another intent's deployment to complete.</li> <li>• In Progress: deployment of the intent is underway.</li> </ul>
Detailed State	This indicates the progress the displayed intent has made toward deployment.
Pipeline User	The login ID of the user who added the intent to the pipeline.
Source Type	The originating intent type: Workload, Maintenance, Backbone, or Fabric
Last Updated Time	The last time the intent was updated.
Source User	The login ID of the user who designed the intent.
Source Name	The name of the intent, as specified in its Name property.
Source Version	The version number of the intent.
Auto Deploy	True or False: indicates whether the auto-deploy option has been enabled.
Status Reason	Additional details for the Status field.

From the Deployment Pipeline page, you can also use the **View** drop-down to display the fabric intent list or view the topology map for the region.

*Figure 12: Choosing another view*



For more information about each of these views, see [The Fabric Intent List view](#) and [The Topology Map view](#) respectively.

From the **Deployment Regions** page, selecting **Fabric Intent List** switches the view to the list of all fabric intents contained within the current region regardless of state of the fabric. From that list, you can perform the following, limited set of fabric operations:

- Open a fabric intent

- Open the deployment pipeline
- Duplicate a fabric intent
- Delete a fabric intent

For more information about each of these operations, see [Fabric intents](#).

## 4.6.1 Viewing the deployment pipeline

### Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** In the main menu, select  **Deployment Regions**. The **Deployment Regions** page opens, showing a graphical representation of regions already created.

**Step 3.** From this page, do either of the following:

- Right-click a region on the map and select **Show Deployment Pipeline** from the contextual menu.
- Click the **Regional List** () button to view the set of regions as a list; then click the **More actions** icon () at the right edge of a region's row and select **Show Deployment Pipeline** from the actions menu.

## 4.6.2 Working with the deployment pipeline

### About this task

From the deployment pipeline, you can manage the set of deployments that are pending for the region. This procedure is a general description of the actions available from the deployment pipeline. For more detailed procedures for deploying fabric intents and workload intents, including special considerations for particular intent types, see [Fabric intents](#), [Workload VPN intents](#), and [Maintenance intents](#).

From the deployment pipeline you can:

- view the design of an intent that is awaiting deployment
- deploy an intent
- remove an intent from the deployment pipeline
- abort a deployment that is underway
- delete a deployed intent from the list

To take any action available for the intents in the deployment queue, do the following:

### Procedure

**Step 1.** Open the deployment pipeline as described in [Viewing the deployment pipeline](#).

**Step 2.** Select an intent from the displayed list and click the **More actions** icon () at the right edge of the row to open the actions menu.

**Step 3.** Do any of the following for a fabric intent:

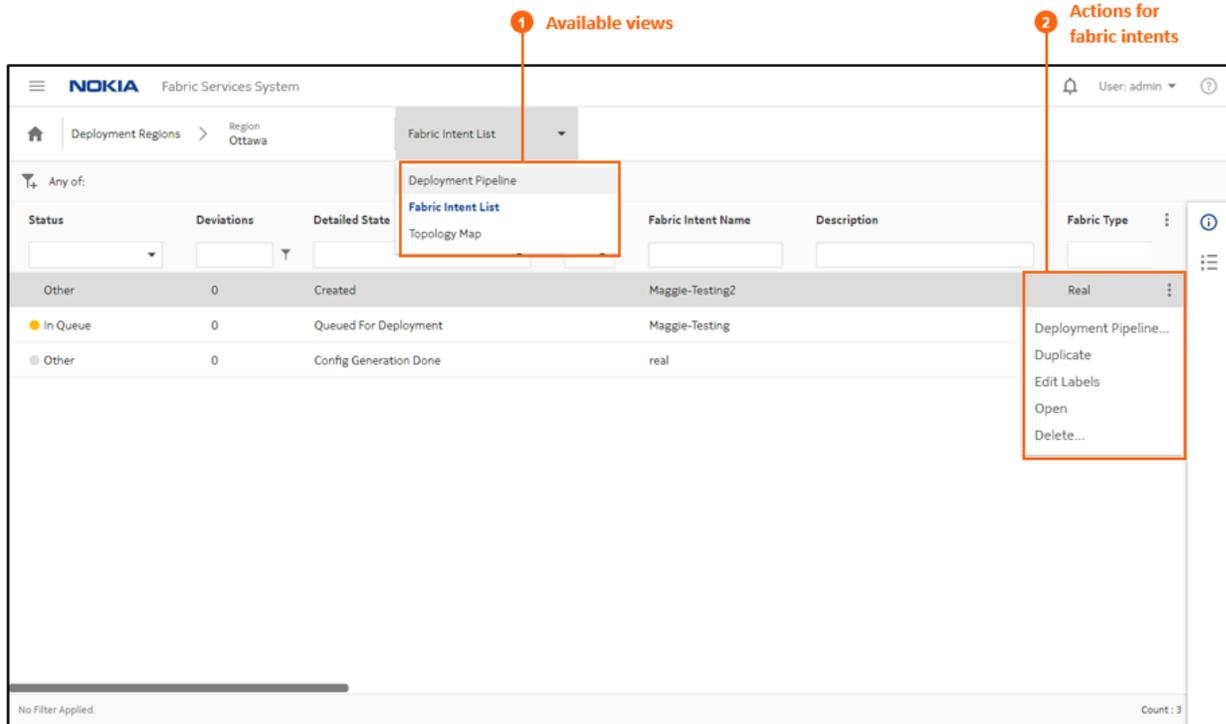
- Select **Open Fabric Design** to leave the deployment queue and view the fabric intent in detail.

- To deploy a flexible leaf/spine fabric intent, select **Deploy**.  
For more information about deploying fabric intents, see [Deploying a fabric intent](#).
- Step 4.** Do any of the following for a workload intent:
- Select **Open Workload Design** to leave the deployment queue and view the workload VPN intent in detail.
  - Select **Deploy** to deploy the workload VPN intent.
- Step 5.** Do any of the following for a maintenance intent:
- Select **Open Maintenance** to leave the deployment queue and view the workload VPN intent in detail.
  - Select **Deploy** to deploy the maintenance intent.
- Step 6.** To remove any intent that has not yet been deployed from the deployment pipeline, do the following:
- a. Select **Remove from pipeline...** in the actions menu. A confirmation dialog displays.
  - b. Click **REMOVE FROM PIPELINE...** in the confirmation dialog. The system removes the selected intent from the deployment pipeline.  
To resume deployment of the intent, you must add it back to the deployment queue from the intent page.
- Step 7.** To cancel the deployment that is underway for any intent, do the following:
- a. Select **Abort Deployment...**
  - b. Click **OK** in the confirmation dialog. The system stops the deployment and rolls back any configured nodes to their preceding state.
- Step 8.** To remove any deployed intent from the list, do the following:
- a. Select **Delete...** A confirmation dialog displays.
  - b. Click **OK** in the confirmation dialog. The system removes the selected intent from the deployment pipeline list.

## 4.7 The Fabric Intent List view

The Deployment Pipeline view includes a drop-down that allows you to switch to alternate views, including the Fabric Intent List view.

Figure 13: The Fabric Intent List view



The Fabric Intent List view shows the list of all fabric intents in the region, along with their current states. From this view you can use selections in the **More actions** (  ) menu to switch back to the **Deployment Pipeline** view, or take actions for individual fabric intents including:

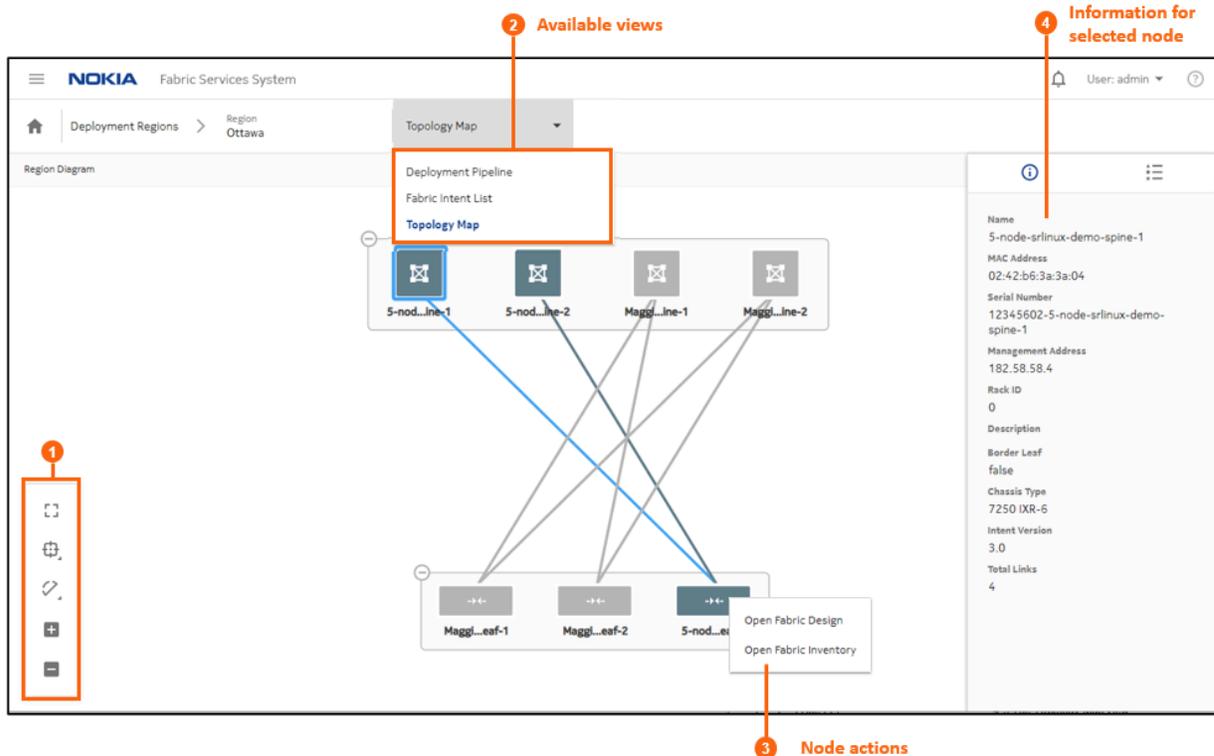
- Duplicate the selected fabric intent
- Edit labels for the selected fabric intent
- Open the fabric intent in the Fabric Design view
- Delete the fabric intent

For more information about each of these actions, see [Fabric intents](#).

## 4.8 The Topology Map view

The **Deployment Pipeline** view includes a drop-down that allows you to switch to alternate views, including the **Topology Map** view.

Figure 14: The Topology Map view



The **Topology Map** view shows the fabrics within the region, and the connections between the leaf, spine, and backbone nodes in those fabrics.

1. Map controls, like those described in [Table 6: Display controls](#), allow you to control the topology display scale, the number of nodes collectively represented by a group icon, and whether links are displayed.
2. The **Views** drop-down menu allows you to switch to the **Deployment Pipeline** or **Fabric Intent List** views.
3. Right-clicking on a node allows you to open the **Fabric Intent Design** or **Fabric Inventory** views, as described in [Fabric intents](#) and [Inventories](#) respectively.
4. Expanding the information panel and clicking on the icon displays information about the currently selected object in the topology view.

## 5 Fabric intents

A fabric is a group of switches that are managed as a single logical unit. A single data center can include many complementary and mutually supporting fabrics.

With the Fabric Services System, you create a new fabric (or plan changes to an existing fabric) by designing and deploying a fabric intent: a detailed plan for a fabric's topology, whose node configuration files are deployed as a single transaction.

Each fabric exists within a region, a logical entity that stores configuration data common to the fabrics within it and that manages the orderly deployment of various types of intents. Regions, and their creation and management, are described in [Deployment regions](#).

After you deploy the fabric intent, the resulting fabric serves as the foundation for additional management templates that you can superimpose to manage the distribution of traffic across the fabric.

### 5.1 Nokia certified templates

The Fabric Services System includes a set of Nokia certified templates you can use to create a fabric intent. Each template is optimized for a particular topology (either flexible leaf/spine or backbone), edge link and T1 layer uplink speeds, and oversubscription settings. Each template displays information to help you select the best template for your planned fabric intent.

When combined with a few initial parameters that you provide, the system uses the network engineering expertise built into the Nokia template to create an optimized fabric topology of nodes and their connection settings.

In the fabric topology descriptions that follow, a table describes some of the default parameters associated with the built-in flexible leaf/spine and backbone templates.

The templates incorporated into the Fabric Services System cannot currently be deleted, altered, or added to.

### 5.2 Fabric topology

The topological design of fabrics for data centers can vary widely. To accommodate diverse fabric structures, the Fabric Services System uses a modular, tiered approach to representing a fabric's hierarchical layers and their interconnections:

- Tier 1: leaf nodes, with edge-link connectivity
- Tier 2: spine nodes, which connect multiple leaf nodes together
- Tier 3: backbone nodes, which connect multiple fabrics together
- Tier 4: regional backbones, connecting backbone nodes together

This flexible, modular approach allows the Fabric Services System to model a wide range of topologies, grouping sets of leaves, spines, and backbone nodes in various ways and distributing connections between and within these groups in whichever way operators deem most efficient.

## 5.2.1 Flexible leaf/spine fabrics

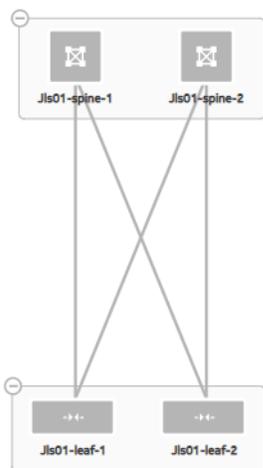
At the base of the fabric hierarchy are the leaf nodes, providing downlink and access ports and connected together by one or more spine nodes.

A single data center fabric could include many such collections of leaf nodes, each with their own set of connecting spine nodes. But regardless of their number, an individual flexible leaf/spine structure is the basic building block of a fabric.

When you create a flexible leaf/spine fabric, you specify the number of leaf nodes and spine nodes; the system then proposes an initial topology in which all leaf nodes are linked to all spine nodes.

**Figure 15: A simple flexible leaf/spine fabric topology** shows a simple representation of a flexible leaf/spine topology, with a pair of leaf nodes (Tier 1) at the bottom of the figure. Each leaf node is linked to both of the spine nodes (Tier 2) at the top of the figure, but not to each other.

Figure 15: A simple flexible leaf/spine fabric topology



It is possible to design a fabric in which these connections are designed more selectively; for example, a particular set of spine nodes in the fabric could be connected only with particular sets of leaf nodes. But initially, the system proposes a fully connected default design that you can then modify if required.

### eBGP peering for PE-CE links: learned route preference

Eventually links are created between leaf nodes and customer equipment; these are PE-CE links. When using eBGP to create such links, multiple paths may present themselves.

#### 5.2.1.1 The Nokia certified flexible leaf/spine template

The Nokia template for a Flexible Leaf Spine fabric uses the following default values:

Table 8: Flexible leaf/spine template defaults

Parameter	Template default value
Spine type	Chassis

Parameter	Template default value
T1 layer uplink speed	G100
Edge link speed	G25
Tier one oversubscription	1:3
Tier two oversubscription	1:1

There are many more parameters that can be set for a fabric intent; this table shows only those that are governed by the flexible leaf/spine template. For more fabric intent parameters, see [Creating a new fabric intent: flexible leaf/spine](#). For the complete list of parameters, see [Appendix: Fabric intent parameters](#).

## 5.2.2 Backbone fabrics

When a data center is composed of two or more flexible leaf/spine fabrics, these fabrics may need to communicate with each other to distribute traffic and share application workloads between them. You may also want to ensure that existing flexible leaf/spine fabrics use a single exit point to communicate with external networks.

The purpose of the backbone fabric is to provide the point of interconnection and shared egress for one or more flexible leaf/spine fabrics.

When you create a backbone fabric intent, you choose the node type to function (singly or more likely in groups) as the backbone, and then select the set of member flexible leaf/spine fabrics that the backbone fabric connects.

For this reason, you must create and deploy the participating flexible leaf/spine fabrics first, and then add the backbone fabric that connects them. Only fabrics in a Deployed state are eligible for selection and attachment to a new backbone fabric intent.

It is always possible to update the design of the backbone fabric intent to add more fully-deployed flexible leaf/spine fabrics as the data center fabric grows.

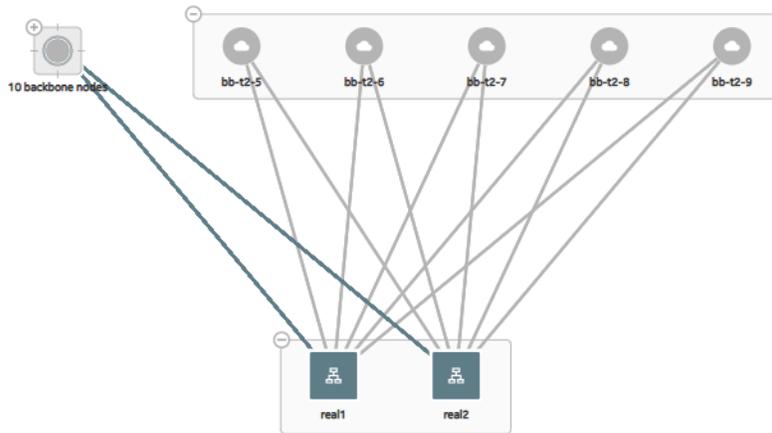
For the purpose of flexibility, the system supports the creation of a backbone fabric with only a single member flexible leaf/spine fabric. Such a design may be warranted if the backbone node provides capabilities that the spine nodes cannot. Or, such a fabric may be a forward-looking design that adds the backbone layer in anticipation of additional flexible leaf/spine fabrics to be added in the future.

When you create a backbone fabric intent and select its member flexible leaf/spine fabrics, there is an implicit change to each of those member flexible leaf/spine fabrics that adds the uplink from the spine node to the new backbone node.

For this reason, when you design a backbone fabric intent, the system automatically creates an updated version of each flexible leaf/spine fabric intent that includes this new uplink. When you deploy the backbone intent, the system first deploys the backbone fabric itself; and then the system deploys each of the updated flexible leaf/spine fabric intents to create the uplink to the new backbone.

**Figure 16: An example of a backbone topology** shows a simple representation of a backbone fabric topology. Fifteen backbone nodes (Tier 3) separated into two groups each connect with two fabrics including their own leaves and spines (Tier 2 and Tier 1). Each flexible leaf/spine fabric is represented here as a single, collective icon.

Figure 16: An example of a backbone topology



### 5.2.2.1 Border leaf

A leaf node can also serve as a backbone node; in this case, it is a "border leaf" node; "leaf" because it can carry EVPN workloads like other leaves, but "border" because unlike other leaf nodes, it communicates with the external network.

Because the node must be capable of carrying EVPN traffic in its role as a leaf node, only EVPN-capable hardware can serve as a border leaf node. This is reflected in the smaller list of hardware options when you select the node type for a border leaf backbone node.

### 5.2.2.2 The Nokia certified backbone template

The Nokia template for a backbone fabric uses the following default values:

Table 9: Backbone template defaults

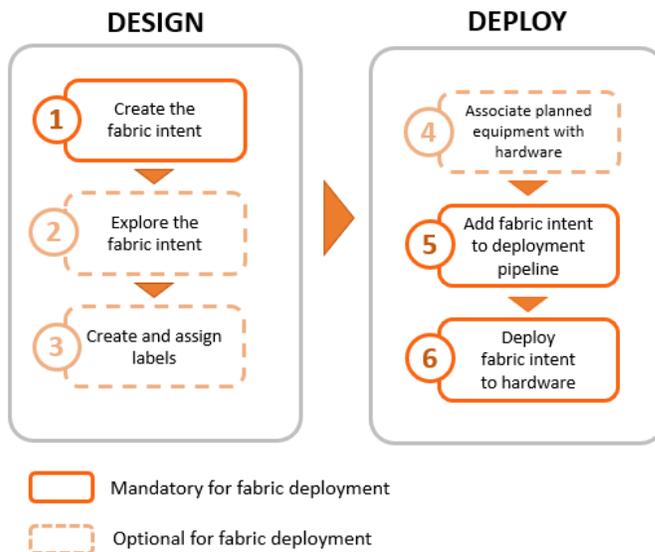
Parameter	Template default value
Node type	7250 IXR-6
Oversubscription	1:3

There are many more parameters that can be set for a fabric intent; this table shows only those that are governed by the backbone template. For more fabric intent parameters, see [Creating a new fabric intent: backbone](#). For the complete list of parameters, see [Appendix: Fabric intent parameters](#).

## 5.3 The design-and-deploy workflow

This section describes the workflow for designing, exploring, and deploying a fabric intent.

Figure 17: Fabric design and deployment workflow



Before you can create any fabric intents, you must:

- create a region (a logical container for fabrics that manages the deployment of intents), as described in [Deployment regions](#).
- ensure that the software catalog includes the SR Linux image you intend for deployment to all nodes in the fabric, as described in [Software and image catalogs](#).

### 5.3.1 Create the fabric intent

Fabric intent creation begins with the selection of a Nokia certified template for the type of fabric type you are creating: flexible leaf/spine or backbone. These templates contain optimized topological and configuration data prepared by Nokia engineers, and can greatly speed up the design of otherwise complicated fabrics and their interconnections.

After you select a template and provide some basic parameters, you save the design and, with one click, generate the detailed fabric intent. The intent includes a fabric topology based on your selections, identifies the types of hardware required for each node, includes detailed configuration code for each node, and creates a wiring plan that technicians can use to connect the hardware together.

The Fabric Services System supports the creation of fabrics containing two types of nodes:

- real nodes, if this is a fabric you intend to deploy onto existing or planned real-world hardware.
- virtual nodes, which the system emulates with software in the Digital Sandbox. Such detailed simulations are useful for exploring, testing, and validating a potential fabric designs. The Digital Sandbox is a separate component that requires its own license. For more information, see [Digital Sandbox](#).

For the procedures to design a fabric intent, see:

- [Creating a new fabric intent: flexible leaf/spine](#)
- [Creating a new fabric intent: backbone](#)

### 5.3.2 Explore the fabric intent

After you create a new fabric intent, you can explore the intent to verify that it appears as expected.

The Fabric Services System includes various tools you can use to explore the resulting fabric topology:

- One is the graphical fabric topology display, whose elements and operation are described in [Viewing a fabric intent](#).
- Another is the fabric-as-code view, that shows the precise configuration data for every element of the fabric. The steps to view a fabric intent or an individual node in code view are described in [Viewing a fabric intent as code](#) and [Viewing the configuration file for a single node](#).

### 5.3.3 Create and assign labels

The Fabric Services System supports an extensive labeling system you can use to tag the elements within a fabric intent. These tags are more than just passive information displays; applied systematically, tags can identify groups of logically related entities that can then be the subject of collective operations in the future. Labels and their application to fabric intents are described in [Labels](#).

### 5.3.4 Associate nodes with hardware

When you are satisfied with the fabric design, you can associate the abstracted, planned nodes in the fabric intent with their real-world counterparts using the hardware serial number. This process is described in [Associating individual planned nodes with real-world hardware](#).

### 5.3.5 Add the fabric intent to the deployment pipeline

The first step in deploying a completed fabric design is to add it to the region's deployment queue as described in [Deploying a fabric intent](#).

Because a large number of intents may require deployment at any one time, the system uses regions to maintain a fixed queue of pending deployments. This deployment pipeline helps the system manage deployments in an orderly sequence and prevents them from interfering with each other.

### 5.3.6 Deploy the fabric intent to hardware

Finally, you can go to the region's deployment pipeline and manually trigger the deployment of the fabric. This downloads the necessary configuration files to each of the participating nodes that you associated previously. This process is also described in [Deploying a fabric intent](#).

## 5.4 Supported hardware

The Fabric Services System currently supports SR Linux hardware in the leaf, spine, backbone, and border leaf roles as shown in [Table 10: Supported hardware and roles](#).

Table 10: Supported hardware and roles

Hardware type	EVPN capable?	Eligible for these roles			
		T1 (Leaf)	T2 (Spine)	T3 (Back-bone)	T3 (Border Leaf)
7250 IXR-6	-	-	Yes	Yes	-
7250 IXR-10	-	-	Yes	Yes	-
7220 IXR-D2	Yes	Yes	-	-	-
7220 IXR-D2L	Yes	Yes	-	-	-
7220 IXR-D3	Yes	Yes	Yes	Yes	Yes
7220 IXR-D3L	Yes	Yes	Yes	Yes	Yes
7220 IXR-H3 <sup>1</sup>	-	-	Yes	Yes	-

1. The system does not configure port speed on 7220 IXR-H3 nodes.

## 5.5 Software and image catalogs

When you create a fabric, you must specify what operating system and version should be present on all of the nodes that participate in the fabric. You select a particular OS software vendor, type, and version number as part of the fabric design, and the system deploys that software to each node as part of the overall fabric deployment.

Currently, the system supports only SR Linux software from Nokia. A set of SR Linux software images are pre-loaded into the system.

Because you select only one SR Linux image file for the fabric intent as a whole, the same version of SR Linux is deployed to all nodes participating in the fabric. If a node is already running a different version of SR Linux, it downloads the new SR Linux software image file and uses it to replace its current version when you deploy the fabric.

For the procedures to view the system's current software catalog, and add a new image to the catalog, see the [Software and image catalogs](#) section in [Application settings](#).

## 5.6 Fabric intents page

The fabric intents page in the Fabric Services System GUI allows you to create and manage individual fabrics.

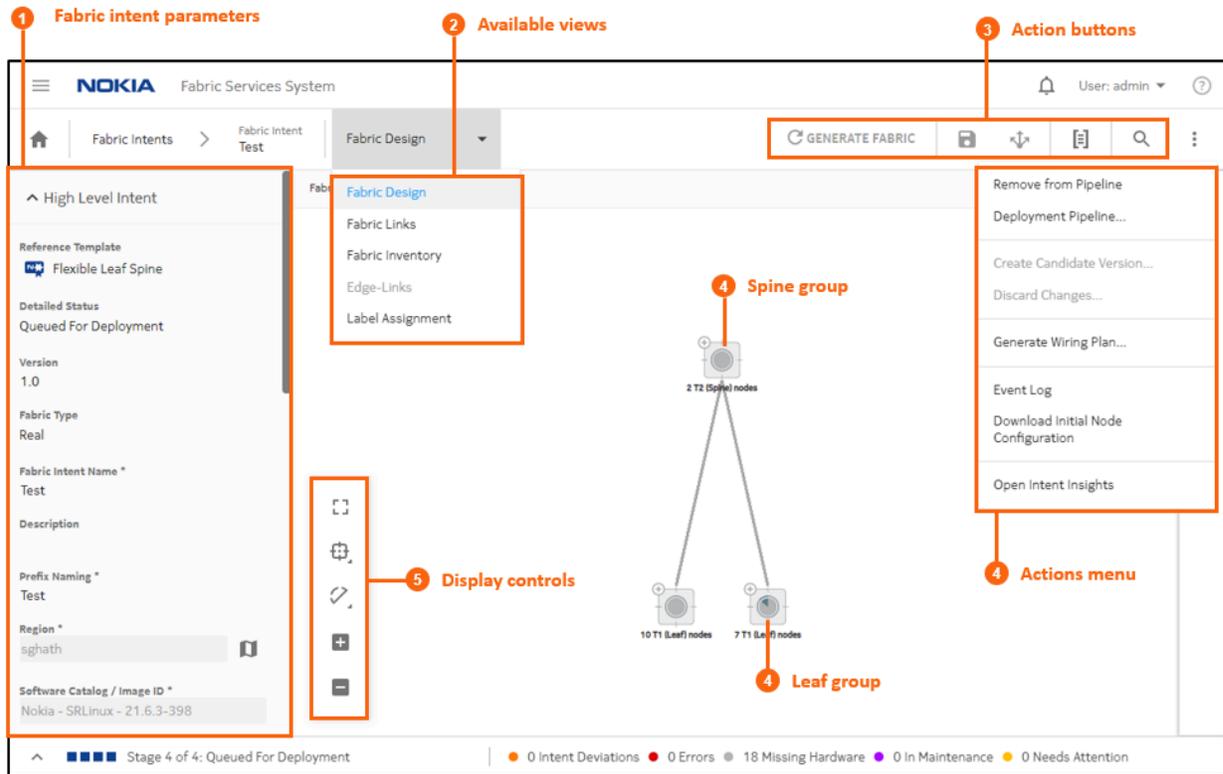


Table 11: Fabric intents page elements

#	Description
1	Fabric intent parameters: when creating a fabric intent, enter basic parameters here. When viewing a fabric intent, configured parameters are displayed here.
2	Available views: several different views available from the this page allows you to interact with different aspects of the fabric intent: <ul style="list-style-type: none"> <li>• <b>Fabric Design:</b> the basic view showing a map of the fabric topology.</li> <li>• <b>Fabric Links:</b> displays a list of links within the fabric topology, and information about each link.</li> <li>• <b>Fabric Inventory:</b> displays a list of all planned nodes within the fabric intent topology, and information about each node.</li> <li>• <b>Edge Links:</b> available only for version 2.0 or greater for a fabric intent, allows you to view and configure the edge links on leaf nodes. From the Edge Links view you can configure LAGs and breakout ports.</li> <li>• <b>Label Assignment:</b> allows you to assign labels to various elements of the fabric intent.</li> </ul>
3	Action buttons: <ul style="list-style-type: none"> <li>• <b>GENERATE FABRIC:</b> use to generate a topology map for a saved fabric intent.</li> <li>• <b>Save</b> (📁): use to save a fabric intent design.</li> <li>• <b>Deploy</b> (📡): use to send a generated fabric intent to the region's deployment queue.</li> </ul>

#	Description
	<ul style="list-style-type: none"> <li>• <b>View fabric as code</b> (📄) : use to view the configuration code for the current fabric; from that view you can compare this version's code against a previous version of the same fabric intent, or view the configuration code for an individual node within the fabric.</li> <li>• <b>Search for node</b> (🔍) : click to search for a node on the topology map; the selected node is highlighted.</li> </ul>
4	<p>Actions menu:</p> <ul style="list-style-type: none"> <li>• <b>Remove from Pipeline</b>: if you have added a fabric intent to the region's deployment pipeline, this action removes it.</li> <li>• <b>Deployment Pipeline...</b>: this action takes you to the Deployment Pipeline page.</li> <li>• <b>Create Candidate Version...</b>: for a deployed fabric intent, this action creates a new version of the same fabric intent.</li> <li>• <b>Discard Changes...</b>: if you are in the midst of creating a new version of a fabric intent, this action discards your changes and resets the page to display the latest version of the intent.</li> <li>• <b>Generate Wiring Plan...</b>: this action creates a downloadable .csv file describing the wiring plan for the current fabric intent.</li> <li>• <b>Event Log</b>: this action opens the event log.</li> <li>• <b>Download Initial Node Configuration</b>: this action creates a downloadable .csv file containing the initial configuration file for each node in the fabric intent.</li> <li>• <b>Open Intent Insights</b>: this action opens the Operational and Health Insights page.</li> <li>• <b>Update Digital Sandbox</b>: if creating a fabric intent destined for the Digital Sandbox, this action updates the Digital Sandbox with the latest configuration data for the fabric intent.</li> <li>• <b>Cancel Digital Sandbox Update</b>: if you have begun an update of the Digital Sandbox, this action cancels that update.</li> </ul>
5	<p><b>Spine group</b>: denotes a collection of nodes acting as the fabric intent's spine. Click the + to expand the cluster and view individual spine nodes. For more information see <a href="#">Groups</a></p>
6	<p><b>Leaf group</b>: denotes a collection of nodes acting as the fabric intent's leaves. Click the + to expand the cluster and view individual leaf nodes.</p>
7	<p><b>Display controls</b>: allows you to configure the topology display. For details see <a href="#">Display controls</a>.</p>

## 5.7 Creating a new fabric intent: flexible leaf/spine

### Prerequisites

Before you create a new flexible leaf/spine fabric intent, ensure you have done the following:

- Created a region to which the fabric can belong; see [Creating a region](#).
- Verified that the intended version of the SR Linux software is available in the system's software catalog, as described in [Software and image catalogs](#).

## About this task

This procedure describes how to create a flexible leaf/spine fabric intent intended for deployment to real hardware.

You can also create a virtual fabric intent for testing in a virtual environment. For the small differences in the fabric design procedure for that purpose, see [Digital Sandbox](#).

A fabric intent normally allows you to select a single hardware type for all leaf nodes within a fabric intent. However, the Fabric Services System supports the inclusion of a second set of leaf hardware with some additional steps. For the procedure to create such a fabric intent, see [Adding a second type of hardware for leaf nodes](#).

## Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** In the main menu, select  **Fabric Intents**.

**Step 3.** Click the **+ CREATE A FABRIC INTENT** button.

Templates are displayed in a grid view by default. To switch to the list view, select  in the template selection screen. Click  to return to the grid view.

**Step 4.** Select the Flexible Leaf Spine template.

**Step 5.** Click the **CREATE** button.

**Step 6.** On the left-side panel, enter or select the basic parameters that define your intended fabric as described in [Table 12: Basic parameters for flexible leaf/spine fabric intent](#)

At the top of the left-side panel, the name of the template you selected in step 4 is displayed as the "Reference Template".

*Table 12: Basic parameters for flexible leaf/spine fabric intent*

Parameter	Description
Fabric Type	Select Real. This creates a fabric intent intended for deployment to real-world hardware. The Digital Sandbox option is used to create a virtual fabric to test and validating prospective designs. This option is described in <a href="#">Digital Sandbox</a> .
Fabric Intent Name	Enter the name that identifies this fabric.
Description	Optionally, add information to describe the fabric intent.
Prefix Naming	Enter a string to be added at the beginning of the name of every node in the fabric intent. The rest of the node name is automatically generated.  For example, enter "A01" here to assign nodes names such as "A01-leaf-1" and "A01-spine-1".  The Prefix Naming string must be unique to each fabric.
Region	Click the drop-down and select an already-created region from the list. This identifies the region that contains the new fabric intent.

Parameter	Description
	To create or edit a region now, click the Regions icon (  ) to navigate to the Regions page. This exits the fabric intent creation procedure. You must begin fabric intent creation again after creating the new region.
Software Catalog/ Image ID	Select the version of SR Linux software to load onto all of the nodes in the new fabric. Currently supported: <ul style="list-style-type: none"> <li>• SR Linux 21.11.2-72</li> <li>• SR Linux 21.6.4-6</li> </ul>
Spine Node Type	Select a supported type of hardware to act as a spine node.
Number of Spine Nodes	Specify the number of devices of the specified type to act as spine nodes.
Leaf Node Type	Select a supported type of hardware to act as a leaf node.   <b>Note:</b> When initially creating a fabric intent, you can only choose one hardware type to be used for all leaf nodes in the fabric. To use a second hardware type for some leaf nodes, you must first deploy the initial fabric intent and then create a new version that includes a pod. The pod represent a second set of leaf nodes within the same fabric which can be of a different hardware type. For more information about configuring fabrics with multiple hardware types for its leaves, see <a href="#">Adding a second type of hardware for leaf nodes</a> .
Number of Leaf Nodes	Specify the number of devices of the specified type to act as leaf nodes.
Labels	This area is disabled when creating the initial fabric intent.
Deploy Trigger Percentage	The minimum percentage of nodes within the fabric intent that must be physically installed before the system allows you to deploy the intent to hardware. For all real fabrics, this value is fixed at 0%. For a full explanation of the impact of this setting, see <a href="#">The Deployment Trigger Percentage setting</a> .
There are more, optional parameters you can use to fine-tune the fabric intent design at this stage. For more information about these parameters, see <a href="#">Appendix: Fabric intent parameters</a> .	

**Step 7.** Click  to save the fabric intent. When you save the fabric intent, the system:

- updates the state of the fabric intent to Created.
- updates the version number of the fabric intent to 1.0.
- enables the  **GENERATE FABRIC** button.

**Step 8.** Click  **GENERATE FABRIC**.

## Expected outcome

The system generates a recommended topology for your fabric based on the template you selected and the parameters you provided.

When generating the topology, the system also generates the various cable connections and the individual node configurations required to support this fabric topology.

During fabric generation, the fabric intent state advances through the following:

- Cable Map in Progress
- Config Generation in Progress
- Config Generation Done

After the generation is complete, the resulting topology displays in the main area of the **Fabric Intents** page.

## What to do next

Now that you have generated the fabric intent, you can proceed to any of the following:

- Optionally, explore the fabric intent in the UI as described in [Viewing a fabric intent](#).
- Optionally, view the fabric intent as code; see [Viewing a fabric intent as code](#).
- Optionally, add labels to objects in the fabric intent; see [Labels](#).
- If possible, associate planned node in the fabric intent with real-world hardware; see [Associating individual planned nodes with real-world hardware](#).
- Begin the deployment process by adding the fabric intent to the region's deployment pipeline. See [Adding a fabric intent to the deployment pipeline](#). From the pipeline you can then deploy the fabric intent.



### Note:

The following configurations are not supported on the initial version of a fabric intent. For the following configurations, you must first need to save and deploy Version 1.0 of the fabric intent, and then create a new version of the intent (as described in [Creating a new version of a fabric intent](#)) in which the following actions are permitted:

- configure LAGs as described in [Creating LAGs](#).
- configure breakout ports as described in [Configuring a breakout port](#).

## 5.8 Adding a second type of hardware for leaf nodes

### About this task

When you create a fabric intent, you must select one type of hardware for all leaf nodes within the fabric.

The Fabric Services System supports fabric designs that use two different types of hardware for leaf nodes; but such fabric intents must be designed in two stages using a feature called Pods.

In the Fabric Services System, a Pod represents a subset of your fabric that uses a different hardware type for its leaf nodes than the original, non-Pod subset. First, you must create and deploy a fabric intent for the non-Pod subset of your fabric. Then you must create a new version of the same fabric intent, add the Pod, and configure this section which can use a different hardware type for its leaf nodes.

The system currently supports adding a single Pod to each fabric intent, which allows for up to two hardware types for leaf nodes within a single fabric.

### Procedure

- Step 1.** Before you begin creating the fabric intent, divide your intended fabric design into two groups of spine and leaf nodes, each of which uses one hardware type for all of its leaf nodes.
- Step 2.** Create the initial fabric intent to represent the first of those groups as described in [Creating a new fabric intent: flexible leaf/spine](#). Select the appropriate number of leaf nodes and their hardware type for the initial group.



**Note:** If you attempt to add a Pod as part of the initial fabric intent, the system displays an error message indicating that Pods are not supported for Version 1.0 of any fabric intent.

- Step 3.** Deploy the initial fabric intent as described in [Deploying a fabric intent](#).
- Step 4.** Create a new version of the same fabric intent using the procedure described in [Creating a new version of a fabric intent](#).
- Step 5.** In this new version, select the appropriate number of spine nodes for the second group.
- Step 6.** Enable the **Add Pod** option. This indicates that you are now configuring the second group within the fabric that uses its own hardware type for leaf nodes.
- Step 7.** Select the appropriate number of leaf nodes and their hardware type for the second group.
- Step 8.** Complete and deploy this new version of the fabric intent.

### Expected outcome

The deployed fabric now includes two separate sets of leaf nodes, each with its own hardware types.



**Note:** After you set the Add Pod option to False (by setting the toggle to the disabled position), you can only make changes to the number of devices and the device type for the previously deployed pod.

## 5.9 Creating a new fabric intent: backbone

### Prerequisites

Before you create a new fabric intent, ensure you have done the following:

- Created a region to which the backbone fabric intent can belong; see [Deployment regions](#).
- Verified that the intended version of the SR Linux software is available in the software catalog, as described in [Software and image catalogs](#).
- Created and deployed the flexible leaf/spine fabrics and that are planned to fall under the new backbone fabric intent. You must select the participating flexible leaf/spine fabrics as part of the backbone fabric creation process.

### About this task

After you create some flexible leaf/spine fabric intents, you can use a backbone intent to help existing fabric intents to communicate with one another, and communicate with an external network through the backbone's uplinks.

This procedure describes how to create a backbone fabric intent intended for deployment to real hardware. You can also create a virtual backbone fabric intent for testing in a virtual environment. For the small differences in the fabric design procedure for that purpose, see [Digital Sandbox](#).

### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** In the main menu, select  **Fabric Intents**.
- Step 3.** Click the **+ CREATE A FABRIC INTENT** button.
- Step 4.** Select the Back Bone template.
- Step 5.** Click the **CREATE** button.

#### Expected outcome

The New Fabric Intents page displays.

At the top of the left-side panel, the name of the template you selected in step 4 displays as the "Reference Template".

- Step 6.** On the left-side panel, enter or select the basic parameters that define your intended fabric as described in [Table 13: Basic parameters for a backbone fabric intent](#):

*Table 13: Basic parameters for a backbone fabric intent*

Parameter	Description
Fabric Type	Choose Real. This creates a fabric intent intended for deployment to real-world hardware.  The Digital Sandbox option is used to create a virtual fabric to test and validating prospective designs. This option is described in <a href="#">Digital Sandbox</a> .
Fabric Intent Name	Enter the name to identify this fabric.
Description	Optionally, add information to describe the fabric intent.
Prefix Naming	Enter a string to be added at the beginning of the name of the new backbone nodes in the fabric intent. The rest of the node name is automatically generated.  For example, enter "A01" here to assign nodes names such as "A01-t3-1".
Region	Click the drop-down and select an already-created region from the list. This identifies the region that contains the new fabric intent.  You must create the backbone fabric intent in the same region as the flexible leaf/spine fabrics with which it is associated. Currently, the system supports only a single region.
Software Catalog/Image ID	Select the version of SR Linux software to load onto all of the nodes in the new fabric.

Parameter	Description
	Currently supported: <ul style="list-style-type: none"> <li>SR Linux 21.11.2-72</li> <li>SR Linux 21.6.4-6</li> </ul>
Is Border Leaf	Enable if this backbone fabric is serving as a "border leaf", able to carry EVPN workloads as well as provide access to an external network.  Disable if this is a standard backbone that is not capable of carrying an EVPN workload.
Node Type	Select a supported type of hardware to act as a backbone node. If this backbone is a Border Leaf, the only eligible selection is 7220 IXR-D3.
Number of Back Bone Nodes	Specify the number of devices of the specified type to act as backbone nodes.
Over Subscription Ratio	Specify the ratio between the bandwidth of the downlink ports and the bandwidth of the uplink ports.
Deploy Trigger Percentage	Specify the minimum percentage of nodes within the fabric intent that must be physically installed before the system allows you to deploy the intent to hardware.  For all real fabrics, this value is fixed at 0%.  For a full explanation of the impact of this setting, see <a href="#">The Deployment Trigger Percentage setting</a> .
Fabric Intents	Select the set of fabric intents that fall under this backbone intent.  Completing this field is described in the next step of this procedure.
Labels	This area is disabled when creating the initial fabric intent.
There are more, optional parameters you can use to fine-tune the fabric intent design at this stage. For more information about all parameters, see <a href="#">Appendix: Fabric intent parameters</a> .	

**Step 7.** Use the **Fabric Intents** field to select the already-deployed flexible leaf/spine fabrics that fall under the new backbone. Add fabric intents by doing the following:

- a. Click the Edit (✎) icon. The Fabric Intents list opens, filtered to show deployed fabrics.
- b. Check the box at the left edge of the row for each fabric you want to connect to the backbone.
- c. Click the **SELECT INTENTS** button. The Fabric Intents page closes, returning you to the Fabric Intent creation page.

**Step 8.** Click  to save a fabric intent. The Generate Fabric button is enabled.

**Step 9.** Click the  **GENERATE FABRIC** button.

## Expected outcome

The system generates a recommended topology for your fabric, based on the fabric template you selected and the parameters you provided. When generating the topology, the system also generates the various cable connections and the individual node configurations required to support this fabric topology. The resulting topology displays in the main area of the Fabric Intents page.

## What to do next

Now that you have created the fabric intent, you can proceed to any of the following:

- Optionally, explore the fabric intent in as described in [Viewing a fabric intent](#).
- Optionally, view the fabric intent as code; see [Viewing a fabric intent as code](#). Optionally, add labels to objects in the fabric intent; see [Labels](#).
- Associate at least some of the planned nodes in the fabric intent with real-world hardware; see [Associating individual planned nodes with real-world hardware](#).

After you associate the planned nodes with hardware, you can add the backbone intent to the region's deployment pipeline. From the pipeline you can then deploy the backbone intent.

## 5.10 Viewing a fabric intent

### About this task

Follow this procedure to view a previously created fabric intent.

### Procedure

**Step 1.** Click the  menu.

**Step 2.** Select  **Fabric Intents**.

**Step 3.** To open a specific fabric intent from the list, do one of the following:

- Double-click the row for that fabric intent.
- Select a row, click the  icon at the right edge of that row, and select **Open** from the displayed action list.

### 5.10.1 Groups

Groups represent multiple nodes in your fabric's topology. When a fabric consists of many nodes, the topology can grow complicated. Groups simplify the display of such complicated structures to make them more comprehensible at a glance.

The maximum number of nodes represented by a single group is a variable you can set with the Clustering Controls display setting (as described in [Table 16: Display controls](#)). The system automatically adds more groups to the display to represent additional nodes beyond this threshold.

You can click the  control for any group to expand it; the node is replaced by the set of individual nodes it represented.

Figure 18: Node group

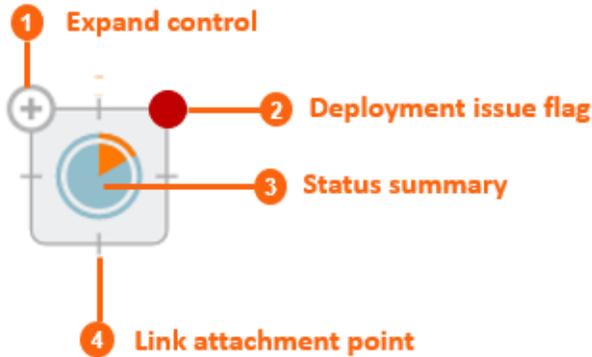


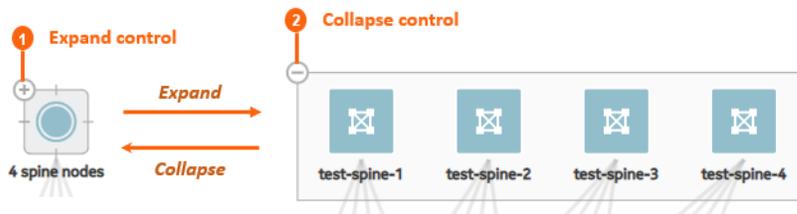
Table 14: Group elements

	Element	Description
1	Expand control	Click this control to expand the group and display its individual nodes.
2	Deployment issue flag	This displays only if there are issues associated with the deployment of configuration data to one or more nodes in the group.
3	Status summary	This is a pie chart that indicates the proportion of nodes in the group that are the subject of alerts, versus those that are not. Errors are represented by the orange segment.
4	Link attachment point	Links to other nodes originate from one of these points at the edge of the group icon.

To expand a group, click the control at the top-left corner of the group.

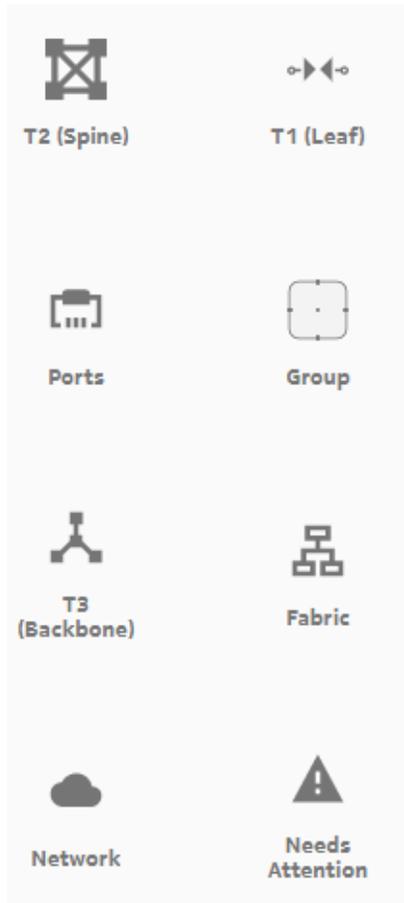
To collapse an expanded group, click the control in the same position.

Figure 19: Expanded group



The symbol for each node indicates its role as spine, leaf, or backbone as shown in [Figure 20: Icons for fabric intents](#).

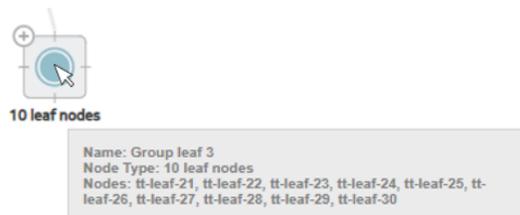
Figure 20: Icons for fabric intents



### 5.10.2 Information displays

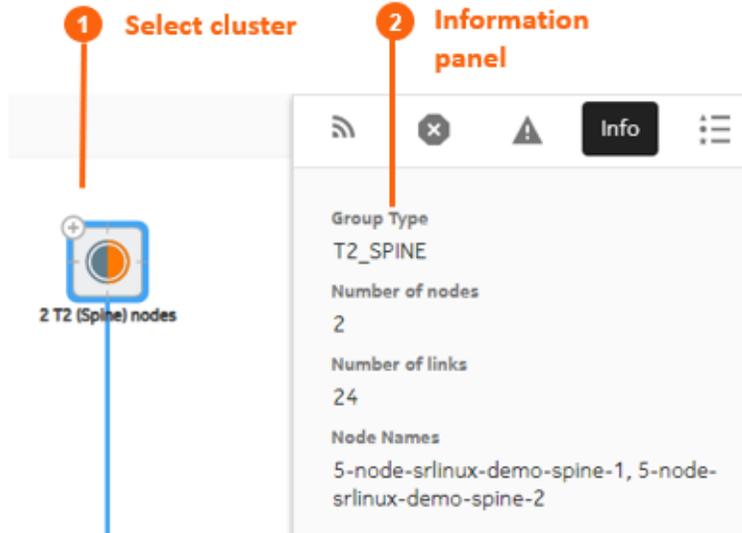
You can hover the cursor over a node or group in the fabric intent Design view to see details about that object.

Figure 21: Hovering over a group



You can also expand the information panel on the right side of the page by clicking . The information panel shows more details about the currently selected object in the fabric topology (either a group, a node, or a link).

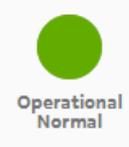
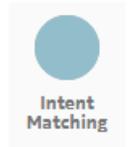
Figure 22: Information panel for a group

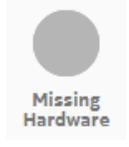
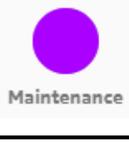


### 5.10.3 Error indicators

The Fabric Services System UI provides a number of cues to draw your attention to problem areas in your fabric intent. Highlights and symbols in the topology display identify groups and nodes that are the subject of one or more errors.

Table 15: Node shading

Element	Description
	Operational Normal: this node or group is in good order and in an operational state.
	Intent Matching: this node or group is configured as described in the original fabric intent.

Element	Description
	Intent Deviation: this node or group has been configured differently than specified in the original fabric intent. This arises when someone outside of the Fabric Services System alters the node configuration after the deployment of your fabric intent.
	Operational Issue: this node or group is not operational or its operation is impaired. Check the Alerts list to identify the operational issue that is affecting the object.
	Missing Hardware: this node or group is identified as "real" in the fabric intent, but no corresponding real-world hardware has been associated with it.
	Maintenance: this node or group is subject to a Maintenance intent (either to replace a node with new matching hardware, or to update the node's software load).

The Status Summary at the bottom of the page shows general information about the fabric intent's progress and status.

Figure 23: Status Summary (expanded)



- Stage: The current fabric intent's progress through its development stages toward deployment, including:
  - Created: A user has selected a template, supplied basic parameters, and elected to proceed with fabric generation.
  - Topology Generated: the system has generated the fabric topology.
  - Configuration Generated: the system has generated the configuration files for individual nodes participating in the fabric, based on the current fabric intent design.
  - Deployment Done: Tracked separately for the Digital Sandbox and real hardware, this indicates that the necessary configurations have been applied to the nodes participating in the fabric.
- Status: These indicators show a count of nodes within the current fabric that are in various states, as described in [Table 15: Node shading](#).

## 5.10.4 Viewing the event log

### About this task

The Fabric Services System maintains a unique event log for every fabric intent, showing significant events in the history of the fabric intent including a time stamp and the outcome of each event.

Figure 24: Event Log for a fabric intent

Event Log			
Timestamp:	2021-03-22T18:53:39Z	State:	Created
Message:	Create Success		
Timestamp:	2021-03-22T18:53:41Z	State:	CableMapInProgress
Message:	Generating Cable Map		
Timestamp:	2021-03-22T18:53:45Z	State:	CableMapDone
Message:	Cable Map Generation Done		
Timestamp:	2021-03-22T18:53:45Z	State:	ConfigGenInProgress
Message:	Generating Config		
Timestamp:	2021-03-22T18:53:45Z	State:	ConfigGenFailed
Message:	ipam precheck fails: Insufficient ISL IP addresses for Links. Request : 288, Total : 128, Available : 62, Unavailable : 2.		

**CLOSE**

To view the event log for a fabric intent:

### Procedure

- Step 1.** Open the fabric intent as described in [Viewing a fabric intent](#).
- Step 2.** From the opened fabric intent, click the  menu at the upper right of the page.
- Step 3.** Select **Event Log** from the action list.
- Step 4.** To exit the event log, click the **CLOSE** button.

## 5.10.5 Display controls

Display controls allow you to manipulate the way objects are displayed within the fabric topology of the Design view.

Figure 25: Display controls on the fabric intents page

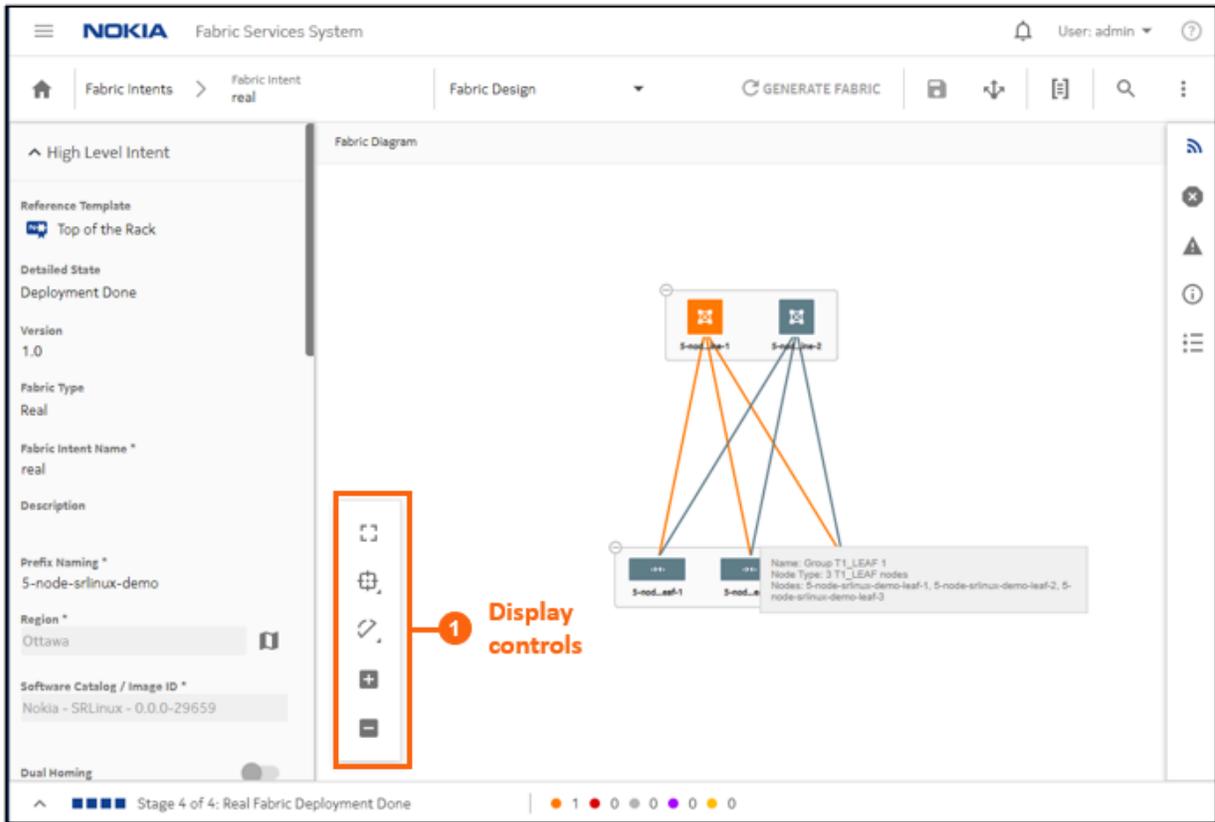


Table 16: Display controls

Icon	Description
	Fit to screen: Adjusts the magnification setting for the topology display so that the entire topology diagram fits inside the current window.
	Clustering controls: <ul style="list-style-type: none"> <li>Enables or disables the use of groups on the map</li> <li>Sets the grouping density (the maximum number of nodes represented by a single group)</li> </ul>
	Adjust links: Allows you to control the way links are displayed in the fabric topology.
	Zoom in, Zoom out: Increases or decreases the magnification for the topology diagram.

## 5.10.6 Viewing a fabric intent as code

### About this task

You can view the configuration code that the system has generated to represent the current fabric. This can be helpful for verifying the design in detail, and possibly revising the fabric design if needed.

### Procedure

**Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).

**Step 2.** Click  to open the code view.

**Step 3.** Optionally do any of the following:

- To save the fabric code in a local file, click  and choose a location and name for the file.
- To copy a portion of the code, select the code and click . The system adds the selection to your clipboard.
- To find a particular string of text within the fabric code, click  and enter the text string. The first instance is highlighted; use the arrows to navigate forward or backward to additional instances, or click ALL to highlight all instances simultaneously.

**Step 4.** To compare the current fabric code against the code for a previous version of the same fabric:

- a. At the end of the "breadcrumb" list at the upper left of the overlay, click the drop-down and select Compare Versions.
- b. Click the **Compare To:** drop-down and select a different version number of the current fabric from the list.
- c. Click the **Expand <number> lines** link to view the full fabric code displayed in two panels, one for each version.
- d. Optionally, repeat step 4.b to select another version of the same fabric against which to compare the current version's code.

**Step 5.** To view the code for an individual node within the fabric:

- a. At the end of the "breadcrumb" list at the upper left of the overlay, click the drop-down and select Fabric Elements.
- b. Click a node in the left column to see its current configuration code.
- c. Click the **Expand <number> lines** link to view the full fabric code displayed in two panels, one for each version.
- d. Optionally, to see the normalized version of the same node's code, click the drop-down at the upper right of the overlay and select Normalized.

**Step 6.** When finished viewing the fabric as code, click the  at the upper right of the overlay to return to the Design view.

## 5.10.7 Viewing the configuration file for a single node

### About this task

You can view the configuration code that the system has generated for each node in the fabric topology. Like viewing the intent as code, this can be helpful for verifying the precise configuration that is planned for the node, and possibly revising the configuration if needed.

To view the current configuration planned for a single node:

### Procedure

**Step 1.** Select the node in the Design view.

**Step 2.** Right-click the node and select **Inspect configuration** from the drop-down list. The configuration file contents display in an overlay window.

**Step 3.** Optionally do any of the following:

- To save the fabric code in a local file, click  and choose a location and name for the file.
- To copy a portion of the code, select the code and click . The system adds the selection to your clipboard.
- To find a particular string of text within the fabric code, click  and enter the text string. The first instance is highlighted; use the arrows to navigate forward or backward to additional instances, or click **ALL** to highlight all instances simultaneously.

**Step 4.** Click the  at the top right of the overlay to return to the Design view.

## 5.10.8 Downloading the initial node configuration

### About this task

You can download a file containing the initial configuration code for all of the nodes participating in the current fabric intent.

To download the configuration file:

### Procedure

**Step 1.** Open the fabric intent as described in [Viewing a fabric intent](#).

**Step 2.** From the opened fabric intent, click the  menu at the upper right of the page.

**Step 3.** Select **Download Initial Node Configuration** from the list. The system immediately downloads the file "initialNodeConfigs" to your Downloads folder.

**Step 4.** To view the configuration, open the initialNodeConfigs file in a text editor.

## 5.10.9 Viewing the fabric inventory

### About this task

You can open a fabric intent's Fabric Inventory view to see a list of all of the nodes that are included in the fabric design. The Inventory view displays extensive details about each node, including its role in the current fabric intent.

This list is a subset of the full inventory; it shows only the nodes participating in the current fabric intent.

From this view you can edit labels for nodes, associate planned nodes with real hardware, inspect an individual node's configuration code, and download a mapping file to help with node association. For more information about the complete inventory and the full set of actions available from any inventory view, see [Inventories](#).

To open the Inventory view for a single fabric intent:

### Procedure

- Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).
- Step 2.** Click the **View** drop-down and select Fabric Inventory.
- Step 3.** To return to the Fabric Design view or open another view, make the corresponding selection from the **View** drop-down.

## 5.10.10 Viewing fabric links

### About this task

You can open a fabric intent's Fabric Links view to see a list of all of the links that are included in the fabric design. The Fabric Links view displays information about each node, including:

- Role
- Endpoint 1 Node, IP Address, and Port
- Endpoint 2 Node, IP Address, and Port
- Any associated labels

From this view you can edit labels on individual links, as described in [Labels](#).

To open the Fabric Links view:

### Procedure

- Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).
- Step 2.** Click the **View** drop-down and select Fabric Links.
- Step 3.** To return to the Fabric Design view or open another view, make the corresponding selection from the **View** drop-down.

## 5.10.11 Viewing edge links

### About this task

You can open a fabric intent's Edge-Links view to see a list of all of the links and LAGs that are included in the fabric design.

To open the Edge-Links view:

### Procedure

- Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#). This opens the fabric intent in the Fabric Design view.
- Step 2.** Click the **View** drop-down and select Edge-Links.

**Expected outcome**

The display switches to the Edge-Links view.

This view displays the available ports in the current fabric intent, as well as any LAGs that have been created within the current fabric intent.

**Step 3.** To view additional details about any edge link:

- a. Select a row and click the more (  ) icon.
- b. Select **Open** from the drop-down menu.

**Step 4.** Click **Close** to return to the Edge-Links view.

To return to the Fabric Design view or open another view, make the corresponding selection from the **View** drop-down.

## 5.10.12 Generating a wiring plan

**About this task**

You can generate a wiring plan to assist technicians when installing and connecting the fabric's planned, supporting hardware.

The wiring plan is in the form of a .csv file that identifies all of the connected pairs of nodes participating in the fabric, their roles (for example leaf or spine), and the ports on each node that serve as endpoints for their interconnections.

To generate a wiring plan:

**Procedure**

**Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#). This opens the fabric intent in the Fabric Design view.

**Step 2.** Click the **More actions** icon (  ) to open the actions menu and select **Generate Wiring Plan...**

**Expected outcome**

The system generates and saves a file with the name <fabric name>\_wiringPlan.csv.

## 5.11 Creating and managing LAGs

You can create Link Aggregation Groups (LAGs) from the Edge Links view. The system supports the creation of LAGs that include multiple ports on a single leaf node, and also multi-home LAGs (MH-LAGs) that include ports on multiple leaf nodes. For MH-LAGS, the ports should be located on leaf nodes within the same fabric.

### 5.11.1 Creating LAGs

**About this task**

You can designate multiple edge links on the nodes managed by the Fabric Services System to be aggregated into a LAG. The system supports the creation of both LACP LAGs and Static LAGs.

**Note:**

LAG configuration is not supported on the initial version of a fabric intent. To configure a LAG, you must first save and deploy Version 1.0 of the fabric intent, and then create a new version of the intent (as described in [Creating a new version of a fabric intent](#)). You can configure LAGs for Version 2.0 and greater of a fabric intent.

**Procedure**

**Step 1.** Choose one of the following:

- If you are configuring a LAG for a fabric intent that is already deployed, begin by creating a new candidate version of the existing fabric intent as described in [Creating a new version of a fabric intent](#). Then go to step 2.
- If you are configuring a LAG for Version 2.0 or greater of a fabric intent that has not yet been deployed, go to step 2.

**Step 2.** Open the fabric intent as described in [Viewing a fabric intent](#).

**Step 3.** Click the **View** drop-down and select Edge-Links.

**Step 4.** Click **+ CREATE A LAG** to begin creating a LAG from the available downlink interfaces.

If **+ CREATE A LAG** is disabled, make sure your fabric intent has not been deployed. If it is deployed, you must create a new candidate version to create LAGs.

**Step 5.** In the LAG Details view, define basic properties for the LAG:

- **Name**
- **LAG Type:** either LACP or Static.
- **MultiHoming:** enable or disable the ability to include ports from different nodes in the same LAG.
- **Single Active:** If you enabled Multihoming, optionally enable **Single Active** to keep a single link active while the other is in standby mode. If **Single Active** is disabled, both links are active.  
After you enable **Single Active**, use the **Preferred Active** drop-down to select the specific link within the LAG to be active.

Enable the **Revertive** option if you want traffic that has been switched to the standby link to return to the preferred active link after the fault is resolved.

- **Lag Speed**



**Note:** The system does not validate that the ports you select as part of the LAG can support the speed you select here. If the configuration is invalid, the system notifies you (and informs you of the reason for failure) when you try to deploy the fabric intent.

**Step 6.** Add the edge link interfaces to constitute the LAG:

a. Click **+ ADD EDGE LINK INTERFACES**.

b. Select any number of edge link interfaces by checking the box at the left edge of the downlink interface's row.



**Note:** If you did not enable MultiHoming, all of the interfaces you select must be located on the same node.

- c. When you have selected the participating edge link interfaces, click **ADD**.

**Expected outcome**

The system returns you to the LAG Details overlay, and the downlink interfaces you selected are now displayed in the Edge Link Interfaces panel.

- Step 7.** Click **SAVE**.



**Note:** You do not need to configure anything in the Local LAG IDs panel. That list updates automatically to show previously configured LAGs.

**Expected outcome**

The system returns you to the Edge Links view, and the LAG you created is now displayed in the list as an additional edge link.

- Step 8.** You can return to the Fabric Design view, or open another view, by making the corresponding selection from the **View** drop-down.

## 5.12 Configuring a breakout port

### About this task

On 7220 IXR-D3 systems, the QSFP28 connector ports (ports 1/3-1/33) can operate in breakout mode. Four breakout ports, each operating at 25G, can be configured on each QSFP28 connector port operating in breakout mode.

In the Fabric Services System, you can enable breakout mode for supporting ports as part of Edge Link Interface configuration for a fabric intent.

When a port is configured in breakout mode, the original port itself can no longer be selected as a sub-interface for a workload intent, nor can it be made part of a LAG. However each of the broken-out sub-ports derived from it can:

- participate in a LAG; however a LAG cannot include a mixture of regular ports and broken-out sub-ports; all of the ports in a LAG must be of one type of the other.
- be selected as a sub-interface for a workload intent.

In lists of ports in the UI, the system displays the broken-out ports as a group of four sub-ports; for example, port 1/3 in breakout mode displays as ports 1/3/1 through 1/3/4.

You can un-configure breakout mode for a port, but only if none of the broken-out ports are currently members of a LAG or functioning as sub-interfaces.



**Note:**

LAG configuration is not supported on the initial version of a fabric intent. To configure a LAG, you must first save and deploy Version 1.0 of the fabric intent, and then create a new version of the intent (as described in [Creating a new version of a fabric intent](#)). You can configure LAGs for Version 2.0 and greater of a fabric intent.

To configure a port in breakout mode, do the following:

### Procedure

- Step 1.** Choose one of the following:

- If you are configuring ports for a fabric intent that is already deployed, begin by creating a new candidate version of the existing fabric intent as described in [Creating a new version of a fabric intent](#). Then go to step 2.
- If you are configuring breakout ports for Version 2.0 or greater of a fabric intent that has not yet been deployed, go to step 2.

**Step 2.** Open the fabric intent as described in [Viewing a fabric intent](#).

**Step 3.** From the **View** drop-down, select **Edge-Links**.

**Step 4.** In the list of edge links, select a port by checking the box at the left edge of its row.

**Step 5.** At the right edge of the row for the selected port, click the **More actions** icon (  ) and select **Open**.

**Step 6.** Click the **Breakout Mode** toggle to enable Breakout Mode for the port.

#### **Expected outcome**

The UI displays the following new properties for the port in read-only mode:

- Num Channels (number of channels), set to 4.
- Channel Speed, set to 25.

**Step 7.** Click **SAVE**.

## 5.13 Modifying a fabric intent

The Fabric Services System provides a number of methods to alter a fabric intent. Options vary depending on whether a fabric intent has been saved, generated, or deployed.

### 5.13.1 Editing a fabric intent

#### **About this task**

You can edit the parameters for a fabric intent that you have previously saved or generated, but not yet deployed. Editing in this way does not create a new version of the fabric intent; the version number is not set, and changes do not increment the version number until you deploy the current fabric intent.

When editing a fabric intent, you cannot change the template that the fabric intent design is based on, but you can alter the description and topology-related values such as the number of leaves and spines. When you begin the edit operation, the system makes clear which fields can be altered and which cannot.

When you deploy a fabric intent, you can no longer edit its parameters. You must create a new version of a deployed fabric intent as described in [Creating a new version of a fabric intent](#), at which point you can save, regenerate, and possibly redeploy the new version.

To edit a saved fabric intent:

#### **Procedure**

**Step 1.** Open the saved fabric intent as described in [Viewing a fabric intent](#).

**Step 2.** Edit parameters in the left panel as required.

**Step 3.** When you are finished editing parameters, click  to save the new configuration.

### 5.13.1.1 Effect of edits on a related backbone intent

As described in [Backbone fabrics](#), a fabric intent may fall within a hierarchical structure under a related backbone intent. When this is the case, any changes you make to a fabric intent must be propagated to the "parent" backbone fabric intent to keep the two descriptions of the same fabric synchronized.

When you modify a flexible leaf/spine fabric intent, the system marks the related backbone intent as "stale". You must resolve this state by manually creating and deploying a new candidate version for that backbone intent.

You do not need to make any changes to the backbone intent during this operation; simply creating a new version and deploying it causes the system to update the backbone intent with the latest version of the participating fabric intents.

### 5.13.2 Discarding changes to a fabric intent

#### Prerequisites

After you create a new version of a fabric intent, but before you deploy it, you can discard the changes you have made. This deletes the new version of the fabric intent entirely, and reverts the display show to the previously deployed version.

If you discard changes for the initial version of a fabric intent that was never deployed, it deletes the fabric intent entirely.

To discard changes for any version of a fabric intent:

#### Procedure

- Step 1.** Open the fabric intent as described in [Viewing a fabric intent](#).
- Step 2.** From the opened fabric intent, click the **More actions** icon (  ) at the upper right of the page.
- Step 3.** Select **Discard Changes**.
- Step 4.** Click **OK**.

#### Expected outcome

The system deletes the current, in-progress candidate version of the fabric intent, and the display reverts to show the last-deployed version.

### 5.13.3 Creating a new version of a fabric intent

#### About this task

After a fabric intent has been deployed, you can still revise its design. This creates a new, incremented version of the previously deployed fabric intent, which you can alter without affecting the currently deployed version. When you are ready, you can deploy the new version of the fabric intent to replace the previous version.

Some actions you perform in the system may require a change to the configurations of one or more fabrics. When this happens, the system may automatically create, save, and generate a new version of the affected fabric intents for you. You can then deploy the new version to apply the necessary configuration changes to the affected nodes.

A new version of a fabric intent keeps the same fabric intent name, but the system assigns an incremented version number, date, and (if necessary) a new user association.

Creating a new fabric intent is also a prerequisite for some operations that are not supported on the initial version of a fabric intent. Currently, the operations that are supported only on Version 2.0 or greater of a fabric intent are:

- configuring LAGs as described in [Creating LAGs](#).
- configuring breakout ports as described in [Configuring a breakout port](#).

To create a new version of an existing fabric intent:

### Procedure

**Step 1.** Open the fabric intent as described in [Viewing a fabric intent](#).

**Step 2.** Click the  menu at the upper right of the page.

**Step 3.** Select **Create Candidate Version**.

#### Expected outcome

The system saves a new version of the current fabric intent.

You can now edit the fabric intent's parameters and save the result without affecting the original version.

**Step 4.** Click **GENERATE FABRIC** to generate a new fabric diagram.

You can now view the fabric intent, and eventually add it to the deployment pipeline and deploy it as you would for a new fabric intent.

At any time before deploying the new version, you can discard this candidate version and revert to the previously deployed version by clicking on the **More actions** icon (  ) in the Fabric Design view and selecting **Discard Changes**.

## 5.13.4 Duplicating a fabric intent

### About this task

In addition to saving new, incremental versions of an existing fabric intent, you can save a completely independent copy of an existing intent. This new duplicate has its own version history that you can continue to develop independently of the original.

To duplicate a fabric intent:

### Procedure

**Step 1.** Open the list of fabric intents as described in [Viewing a fabric intent](#).

**Step 2.** Click the **More actions** icon (  ) at the right end of the row.

**Step 3.** Select **Duplicate** from the action list.

#### Expected outcome

The system creates a new, independent copy (not a new version) of the selected fabric intent.

## 5.14 Deleting a fabric intent

### About this task

You can delete a fabric intent. This removes the fabric intent and all of its versions from the Fabric Services System database.

You cannot delete a flexible leaf/spine fabric intent that falls under a deployed backbone fabric intent.

Deleting a backbone fabric intent automatically creates and deploys new versions of all participating fabric intents, which remove all references to the former backbone.

The system does not delete the management IP on real nodes when you delete a fabric intent.

To delete a copy of a fabric intent:

### Procedure

- Step 1.** Open the list of fabric intents as described in [Viewing a fabric intent](#).
- Step 2.** Select the row corresponding to the fabric intent you want to delete.
- Step 3.** Click the **More actions** (  ) at the far-right edge of the row.
- Step 4.** Select **Delete...** from the action list and click **OK** in the confirmation dialog.

## 5.15 Associating planned nodes with real hardware

### Prerequisites

Before you can associate a fabric's planned nodes with hardware:

- you must have created and saved the fabric intent, and generated the fabric topology.
- the inventory must include entries for the real-world nodes that correspond to the planned nodes in the fabric intent. The inventory can gain entries for nodes in two ways:
  - when the node hardware is installed, the system automatically discovers the hardware as described in [Node discovery](#).
  - before the node hardware is available, you can pre-load the inventory with information about the anticipated hardware as described in [Inventories](#).

### About this task

You can use the Inventory view to associate planned nodes in the fabric intent with real-world hardware. The system can only download the necessary configuration data to a node within the fabric intent after this association has been made.

This procedure describes how to manually associate each node with hardware, one node at a time.

However, the system also supports uploading a mapping file containing information for multiple nodes

For the procedure for uploading a mapping file, and a more complete description of the available options for associating nodes, see [Inventories](#).

### Procedure

- Step 1.** Open the fabric intent as described in [Viewing a fabric intent](#).
- Step 2.** In the View drop-down, select **Fabric Inventory**.
- Step 3.** Select a row corresponding to one of the planned nodes and do the following:

- a. Click in the **Serial Number** column for that row.
- b. Enter the serial number for the real-world node that corresponds to this planned node.
- c. Wait for the real-world node to adopt a Ready state.

## 5.16 Deploying a fabric intent

Deploying a fabric intent creates a functioning instance of the fabric on any associated, real hardware.

Before you can deploy your fabric, you must have saved the fabric intent and generated its configuration.

Deploying the fabric involves two procedures:

1. [From the fabric intent page, add the fabric intent to the region's deployment pipeline.](#)
2. [From the deployment pipeline, select the fabric intent and select Deploy from the actions menu.](#)

### 5.16.1 Adding a fabric intent to the deployment pipeline

#### Procedure

**Step 1.** Open the list of fabric intents as described in [Viewing a fabric intent](#).

**Step 2.** Click .

**Step 3.** Click the **ADD TO PIPELINE** button.

#### Expected outcome

The system adds the fabric intent to the deployment queue for the region, and updates the fabric intent's Status to "Queued for Deployment."

#### Expected outcome

After you add a fabric intent to the deployment pipeline, it remains in the pipeline until you tell the system to proceed with the deployment. This can be useful if the hardware for the fabric is not yet in place; you can "park" the fabric intent in the deployment queue until the hardware is present, and only then trigger deployment.

#### What to do next

When you are ready to proceed with deployment, go to [Deploying a fabric intent from the deployment pipeline](#).

### 5.16.2 Deploying a fabric intent from the deployment pipeline

#### About this task

When you are ready to deploy the fabric intent to hardware, do the following:

#### Procedure

**Step 1.** Click the  menu.

**Step 2.** Select  **Fabric Intents**.

- Step 3.** Select the row corresponding to the fabric intent you want to deploy.
- Step 4.** Click the  icon at the right edge of the row and select **Deployment Pipeline...** from the actions list.
- Step 5.** Find your fabric intent in the deployment pipeline list.
- Step 6.** At the right edge of the row click the **More actions** icon (  ).
- Step 7.** Select **Deploy** from the actions list.

#### Expected outcome

- For a flexible leaf/spine fabric intent, this causes the system to deploy the fabric.
- For a backbone fabric intent, this causes the system to:
  - a. deploy the backbone intent that connects the member flexible leaf/spine fabrics.
  - b. deploy, serially, new versions of each associated flexible leaf/spine fabric with the necessary information to connect them to the pending backbone intent.  
For more information about the special behavior when deploying a backbone intent, see [Deploying backbone fabric intents](#).

#### Expected outcome

When deployment is complete, the fabric's Status advances to Deployed.

The timing of the deployment depends on the fabric intent's Deployment Trigger Percentage; see [The Deployment Trigger Percentage setting](#).

If any issues arise during deployment, they appear as errors in the error panel at the right of the page. The system also highlights them as deployment issues in the status bar, and by adding a red circle to the nodes and groups affected by the deployment error.

### 5.16.3 The Deployment Trigger Percentage setting

When you create a fabric intent, you select a Deployment Trigger Percentage value. This value represents the percentage of nodes participating in the fabric intent that must be in a Ready state before the system begins deploying the fabric intent.

The Fabric Services System currently supports two values for the Deployment Trigger Percentage: 0% and 100%:

- All Real fabrics use the 0% value, and this cannot be changed.
- All Digital Sandbox fabric intents use the 100% value, and this cannot be changed.

Depending upon the fabric's Deployment Trigger Percentage value, the system does one of the following after you select **Deploy** in the Deployment Pipeline action menu:

- If the value is set to 0%, the system immediately begins deploying the configuration to any available Ready nodes.
  - The system deploys the necessary configuration to each Ready node. The system waits until the configuration is complete on one node before proceeding to the next node.
  - When the system runs out of Ready nodes to configure, including if there were no Ready nodes to begin with, it suspends the deployment process.  
The State of the fabric intent advances to Deployment Done.

- Deployment of the fabric intent does not resume, even if additional nodes reach a Ready state, until you:
  - open the Deployment Pipeline
  - select the fabric intent
  - select **Deploy**At that point the system resumes deployment until it again runs out of Ready nodes, or the entire fabric intent is deployed.
- If any node configuration fails, the system rolls back the entire fabric deployment. Each node that it reconfigured is restored to its previous state, as captured in the rollback save.
- if the value is set to 100%, the system waits until all nodes in the fabric intent are in a Ready state before beginning deployment.
  - After all nodes are Ready, the system deploys the necessary configuration to each node. The system waits until the configuration is complete on one node before proceeding to the next node.
  - The system saves a rollback configuration before deploying a configuration, in case it needs to roll back the deployment on the node.
  - If deployment of the entire fabric completes successfully, the system sets the State of the fabric to Deployment Done.
  - If any node configuration fails, the system rolls back the entire fabric deployment, one node at a time. Each node that it reconfigured is restored to its previous state, as captured in the rollback save.

### 5.16.4 Deploying backbone fabric intents

A backbone fabric connects multiple leaf-spine fabrics together. When you create a backbone fabric intent, you must identify each of the member leaf-spine fabrics that connects through the new backbone.

When you save the newly designed backbone intent, the system simultaneously creates a new version of each of the member flexible leaf/spine fabrics. These new fabric intents are updated to include the new, mandatory connection to the planned backbone fabric.

When you deploy the backbone fabric from the deployment queue, the system performs several operations in sequence:

1. It deploys the backbone fabric intent itself.
2. It deploys the revised version of one of the member flexible leaf/spine fabrics, now updated to include the new connection to the just-deployed backbone fabric.
3. It repeats step 2 until it has deployed new versions of all of the flexible leaf/spine fabrics that are connected to the new backbone fabric.

## 5.17 Deviations

After you deploy a fabric intent, various circumstances may alter the node configurations from those expressed in your fabric intent's configuration files. The system continues to monitor all nodes in the fabric after deployment, and flags any configuration change from your last-deployed fabric design as a "deviation".

Deviations are highlighted in several places in the UI:

- a node whose configuration includes at least one deviation is highlighted in orange on the fabric diagram.
- a group that contains one or more nodes with deviations displays a status summary pie chart on the group icon; this chart is shaded orange in proportion to the number of nodes in the group that are subject to alerts, including deviations.
- the status summary at the bottom of the Fabric Design view shows the number of deviations detected throughout the fabric.
- clicking the deviations display in the status summary opens the Alerts log, which lists individual deviations; and like any alert, you can double-click it to view additional details.

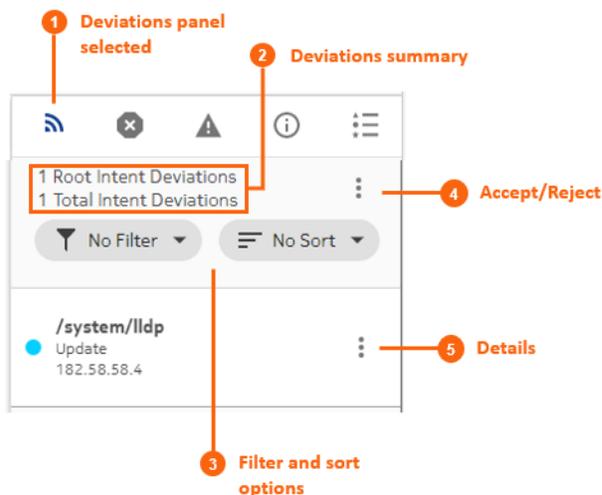
If you attempt to deploy a new version of a fabric intent when there are deviations in the currently-deployed version of that fabric, the system displays an error and prevents your deployment. The error message redirects you to the Alerts panel for fabric intent, where a list of deviations is displayed and you can accept or reject each deviation listed there.

### 5.17.1 Viewing deviations

#### About this task

The Deviations (🔔) panel displays a list of all variations between the fabric's last deployed configuration and its current configuration.

Figure 26: Deviations panel



For each item in the Deviations panel, you can view more information. You can also use the panel to either accept or reject deviations as described in [Accepting or rejecting deviations](#).

To view deviations using the Deviations panel:

#### Procedure

**Step 1.** Click 🔔 on the right side of the fabric intents page.

- Step 2.** Optionally, use the **Filter** drop-down and select a filter type:
- No Filter
  - Delete
  - Update
- Step 3.** Optionally, use the **Sort** drop-down to select one of the following sort methods:
- No Sort
  - Ascending
  - Descending
- Step 4.** Optionally, click the **More actions** icon (  ) to the right of any deviation and select **Details** to view more information about that particular deviation.

## 5.17.2 Accepting or rejecting deviations

### About this task

After the system notifies you of deviations that have been configured on the nodes within a fabric, you have the option to either accept or reject the deviations.

- If you accept the deviations, the system creates a new candidate version of the fabric intent that includes the affected node. The new candidate version includes an updated configuration that incorporates the deviation. The system automatically deploys the updated fabric intent.
- If you reject the deviations, the system creates a new candidate version of the previous fabric intent and automatically deploys it, overwriting the deviation and restoring each node to its pre-deviation configuration.

You can only accept or reject deviations while a fabric intent has been fully deployed (indicated by the "Deployment Done" state).

You can only accept or reject a set of deviations once for a candidate version of a fabric intent. If additional deviations arise before you have deployed the current version, and you attempt to accept or reject any or all of them, the system displays an error message indicating that you cannot do so because the current candidate version of the fabric intent is not in a Deployment Done state.

After you accept or reject the deviations for a fabric intent, the system automatically deploys the resulting new configuration of the fabric intent.



**Note:** The system will not handle any new deviations while deployment is in progress for previous deviations.

To accept or reject deviations to a deployed fabric intent:

### Procedure

- Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).
- Step 2.** Click the Deviations icon (  ) to open the Deviations panel.
- Step 3.** Click the **More actions** icon (  ) to the upper right of the Deviations panel, and select **Accept/Reject** from the displayed menu. The Accept/Reject Watches overlay displays.
- Step 4.** In the displayed list of deviations, select one or more deviations by checking the box to the left of each row. This should be a set of deviations that you want to either accept or reject as a whole.

Alternatively, you can click **Accept All** or **Reject All** to accept or reject the entire list of deviations.

**Step 5.** To accept the selected deviations:

- a. Click the **Accept** button at the bottom of the overlay. The system displays a confirmation dialog indicating that the selected items will be accepted, and all other items in the list will be rejected by implication.
- b. Click OK to confirm the action. The confirmation dialog closes.
- c. Click the **X** at the upper right of the overlay to return to the Design view. Note that the system creates a new candidate version of the fabric intent design.

**Step 6.** To reject the selected deviations:

- a. Click the **Reject** button at the bottom of the overlay. The system displays a confirmation dialog indicating that the selected items will be rejected, and all other items in the list will be accepted by implication.
- b. Click OK to confirm the action. The confirmation dialog closes.
- c. Click the **X** at the upper right of the overlay to return to the Design view. Note that the system creates a new candidate version of the fabric intent.

### Expected outcome

After you accept or reject the deviations for a fabric intent, the system automatically deploys the resulting new configuration of the fabric intent.

## 6 Workload VPN intents

A workload VPN intent defines a set of VPN resources that can be made available to serve a particular source of demand for a data center's traffic, processing, and storage capacity.

Just as with fabric intents, the Fabric Services System uses the concept of an "intent" to describe a set of configuration data that affects multiple nodes, and are deployed together.

When deploying a workload VPN intent:

- the system deploys all of the new configurations for participating nodes as a single transaction. This ensures that the entire deployment succeeds completely, or else fails completely and cleanly, leaving nodes in their pre-deployment states to facilitate another deployment attempt
- the fabric intent remains in the same intent version

### 6.1 Role of a workload intent

A workload intent represents a specific subset of fabric resources (one or more fabrics, subnets, and sub-interfaces). After it has been defined, a workload intent can be allocated to a single source of demand upon the participating fabric or fabrics. For example, all of the traffic from a single customer, or tenant, can be directed to only those resources encompassed by a particular workload intent.

The workload intent is embodied in a set of node configuration files that, when the intent is deployed, cause the participating nodes to behave as described in the intent.

### 6.2 Elements of a workload intent

A workload VPN intent must be associated with specific fabrics within a deployment region (or "region"). The region manages the deployment pipeline, which organizes the sequential deployment of configuration data to the various nodes participating in the intent. For more information about regions, see [Deployment regions](#).

To represent the set of resources available to a workload, a workload intent uses a hierarchical structure consisting of:

- **Fabrics:** a workload intent can span one or more fabrics that you have already defined in the system.
- **Subnets:** each subnet defines a subset of the participating fabrics. Subnets can be either bridged or routed.

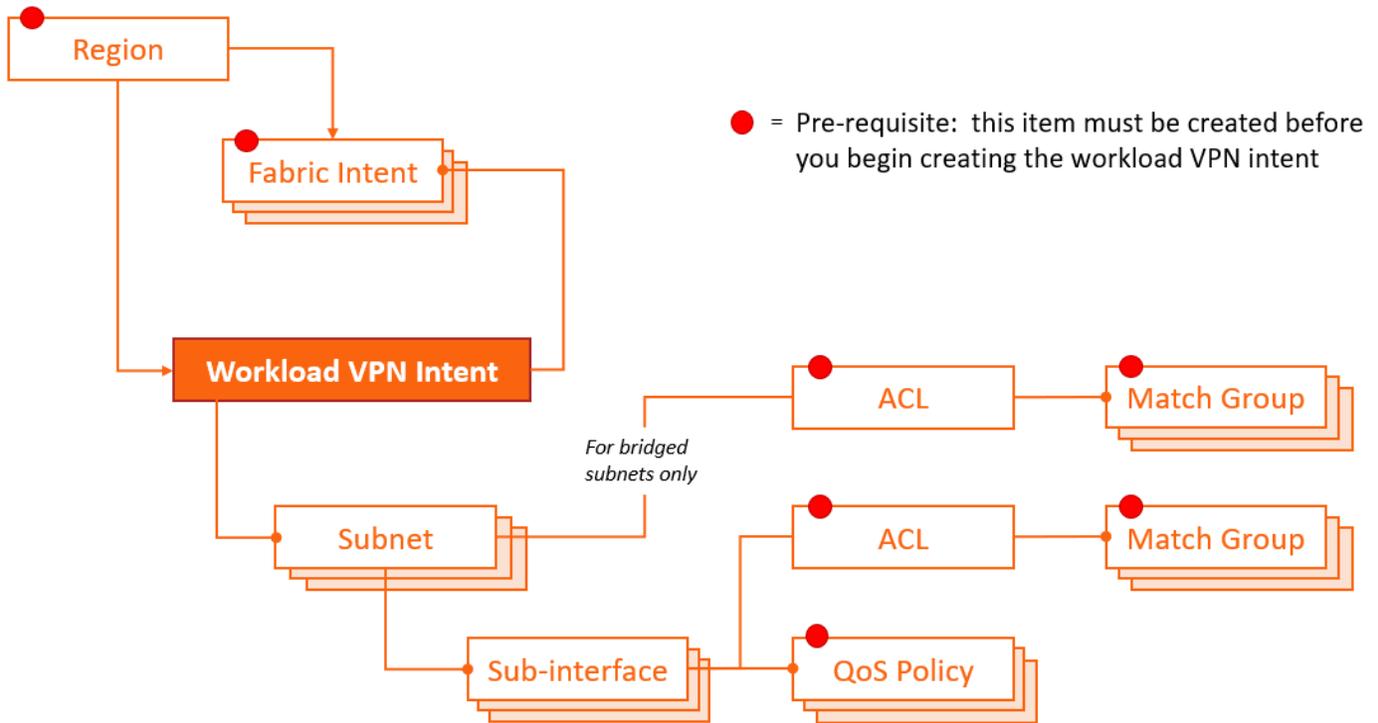
A bridged subnet with one or more IP addresses configured can be assigned an Access Control Lists (ACL); a set of rules that permit, or reject, any traffic that matches those IP addresses. These rules are applied on the IRB interface. The set of IP addresses is stored as a Match Group; the ACL then assembles a set of match groups, adds further defining properties, and indicates whether to accept or reject the matching packets.

- **Sub-interfaces:** each sub-interface identifies a set of edge link ports and LAGs within a particular subnet, and across which the workload traffic can be bridged or routed.

A sub-interface can be associated with:

- Access Control Lists (ACL): each sub-interface can be assigned a set of rules that permit, or reject, any traffic that matches those IP addresses. The set of IP addresses is stored as a Match Group; the ACL then assembles a set of match groups, adds further defining properties, and indicates whether to accept or reject the matching packets.
- QoS profiles: each sub-interface can be assigned a set of rules for prioritizing traffic.

Figure 27: Elements of a workload VPN intent



### 6.2.1 Fabrics

Every workload intent allocates resources that belong to one or more fabric intents that already exist in the system.

As part of workload intent design, you must select the participating fabric intent or intents. All traffic managed by the workload intent is confined to these participating fabric intents (and, more specifically, only to the subnets and sub-interfaces you select within those fabric intents).

A single fabric intent can support multiple workload intents.

### 6.2.2 Subnets

A sub-network, or subnet, is a set of one or more nodes on a selected fabric.

When creating a workload intent, you identify each subnet as either bridged or routed.

- Bridged subnets can be assigned to a network-instance of type MAC-VRF. They are associated with a MAC-VRF network instance, and allow for configuration of bridge table and VLAN ingress/egress mapping.

When adding a bridged subnet to a workload VPN intent, you associate an IP Anycast Gateway, and can associate an ACL.

With Gateway IP, a MAC-VRF is attached to the default or ip-vrf network-instance by a single Integrated Routing and Bridging (IRB) interface. This allows routing between different MAC-VRF instances of a particular tenant or group of servers.

- Routed subnets can be assigned to a network-instance of type mgmt, default, or ip-vrf. They allow for configuration of IPv4 and IPv6 settings.

### 6.2.3 Sub-interfaces

For each subnet, you identify the set of sub-interfaces over which tenant traffic can travel. These sub-interfaces can be physical or logical ports, including Link Aggregation Groups (LAGs).

Each LAG you select as a sub-interface must have been configured previously within the system as part of the Fabric Intent design, and the interfaces participating in a LAG can be located on a single node, or on multiple nodes (constituting a multi-home LAG or MH-LAG).

For each sub-interface, you can define Quality of Service (QoS) profiles that prioritize traffic, and an Access Control List that identifies the packets to be accepted or rejected.

### 6.2.4 QoS profiles

The Fabric Services System supports Quality of Service (QoS) policies for assigning traffic to forwarding classes and remarking traffic at egress before it leaves a router. DSCP classifier policies map incoming packets to the appropriate forwarding classes, and DSCP rewrite-rule policies mark outgoing packets with an appropriate DSCP value based on the forwarding class.

After you have defined one or more QoS policies, you can assign them to workload intents to help manage the traffic associated with that workload.

For the procedure to create a QoS profile, see [Creating a QoS profile](#).

For the steps to apply a QoS profile within a workload intent, see the "Create a sub-interface" section of the procedure [Creating a workload VPN intent](#).

### 6.2.5 ACL profiles

An Access Control List (ACL) profile defines a set of packet types that should be either accepted or rejected. When assigned to a subnet (for bridged subnets) or sub-interface (routed or bridged), the ACL defines the traffic that is permitted on those workload intent resources.

An ACL profile is assembled from one or more match groups. Each match group defines a particular set of properties that could be possessed by packets. The ACL then uses these match groups to specify a set of packet properties, and then includes an instruction to either accept or reject packet that conform to those properties.

Match groups thereby provide an easy way to define a set of IP addresses once, and then re-use that information in multiple ACLs.

For the procedure to create a match group, see [Creating a match group](#).

For the procedure to create an ACL profile, see [Creating an ACL profile](#).

For the steps to apply an ACL profile within a workload intent, see the procedures [Adding subnets to the workload intent](#) (for bridged subnets) and [Adding sub-interfaces to the workload intent](#) (for routed subnets).

## 6.3 Viewing a workload intent

### Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** From the menu, select  **Workload VPN Intents**.

#### Expected outcome

The **Workload VPN Intents** page displays, showing a list of already-created workload intents in the system.

**Step 3.** To open a specific workload intent from the list, do one of the following:

- a. Double-click the row for that workload intent.
- b. Select a row, click the  icon at the right edge of that row, and select **Open** from the displayed action list.

### 6.3.1 Viewing a workload intent

Figure 28: Workload intent Design view

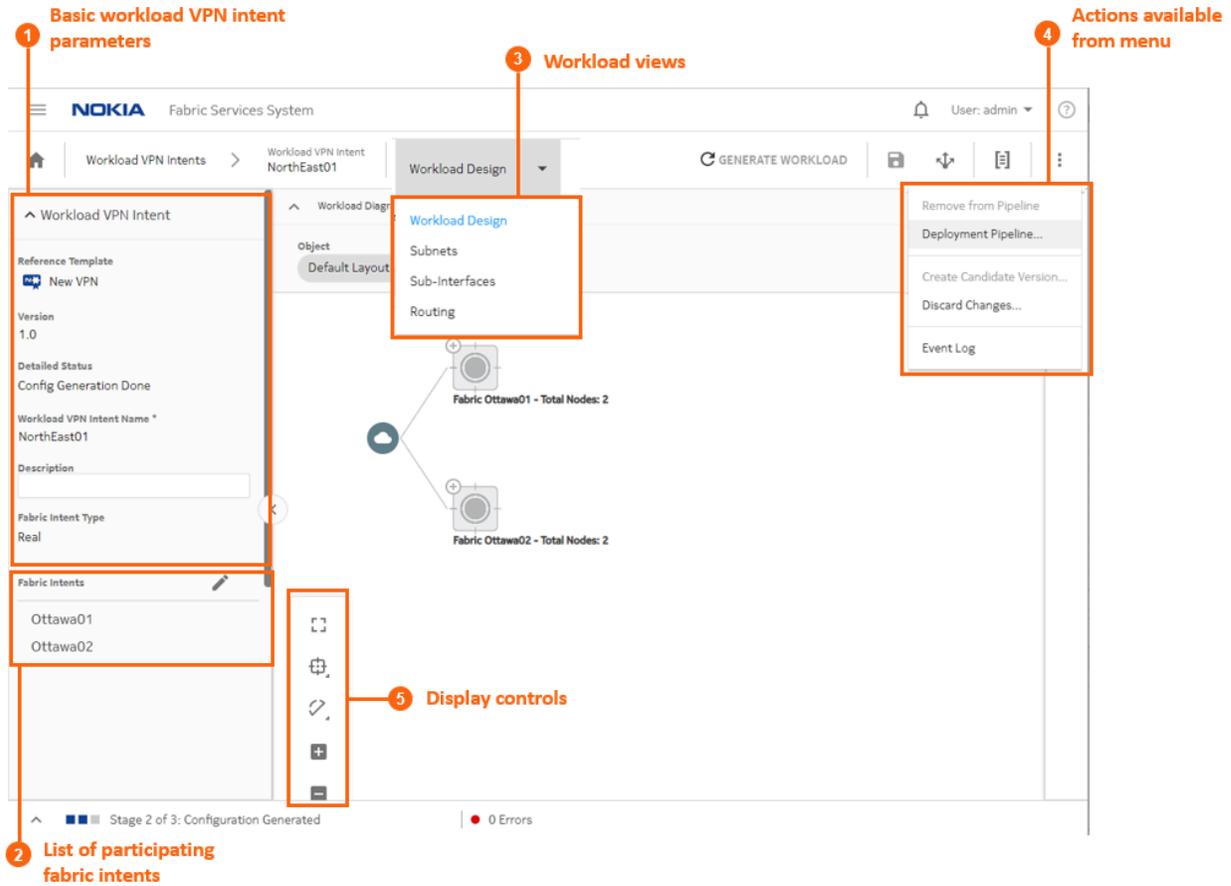
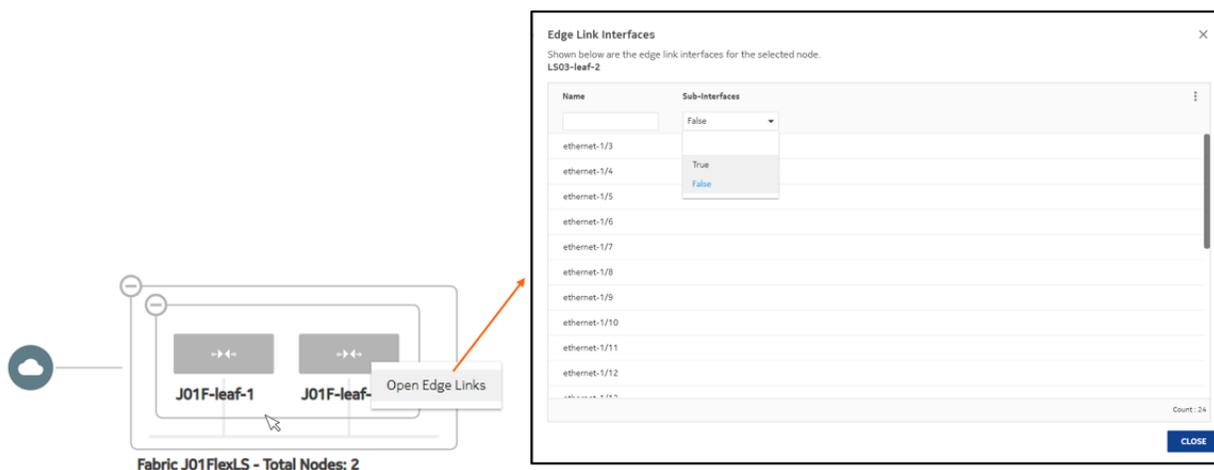


Table 17: Elements of workload VPN intent Design view

	Property	Description
1	Basic workload VPN intent parameters	Basic parameters include the VPN version number and state, as well as its name and a description.
2	List of participating fabric intents	This list shows the set of fabric intents over which the workload intent can distribute traffic.
3	Workload views	The View drop-down provides access to the different views necessary to see and design a workload VPN intent, its subnets, and its sub-interfaces. The Routing view allows you to configure BGP settings.

	Property	Description
4	Actions available from menu	From the actions menu you can manage deployment, create a new version of the intent, or view the intent's event log.
5	Display controls	These controls allow you to modify the way the system displays the workload intent.

Figure 29: Workload VPN intent map objects



On the map display for workload VPN intents, you can expand backbone clusters to show the participating backbone nodes, and leaf clusters to show participating leaf nodes.

You can also right-click a leaf node and select **Open Edge Links** to open an overlay listing all of the edge links for the selected node.

### 6.3.2 Viewing a workload intent as code

#### About this task

You can view the detailed configuration code that the system has generated to represent the current workload intent. This can be helpful for verifying the design in detail, and possibly revising the workload intent design if needed.

#### Procedure

- Step 1.** Open a workload intent as described in [Viewing a workload intent](#).
- Step 2.** Click  to open the code view.
- Step 3.** Optionally, do any of the following:
  - To save the workload intent code in a local file, click  and choose a location and name for the file.

- To copy a portion of the code, select the code and click . The system adds the selection to your clipboard.
- To find a particular string of text within the workload intent code, click  and enter the text string. The first instance is highlighted; use the arrows to navigate forward or backward to additional instances, or click **ALL** to highlight all instances simultaneously.

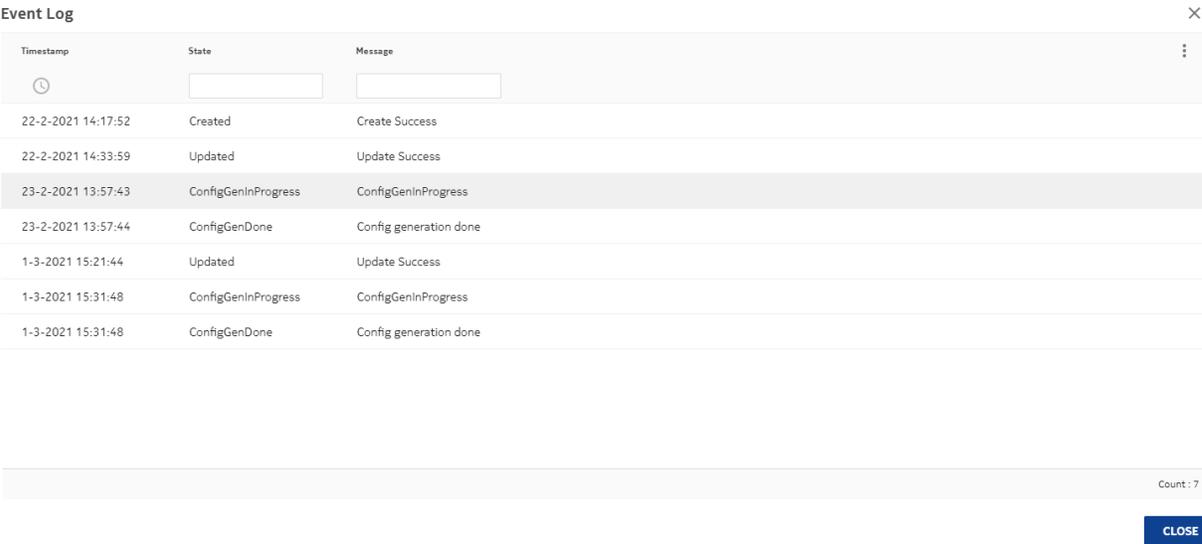
**Step 4.** Click  at the upper right of the overlay to return to the **Workload Design** view.

### 6.3.3 Viewing the event log

#### About this task

The event log retains a detailed history of all registered events that have occurred with respect to the current workload intent. This can be useful when troubleshooting workload intent issues, or just to verify that the workload history matches expectations.

*Figure 30: The workload VPN intent event log*



Timestamp	State	Message
22-2-2021 14:17:52	Created	Create Success
22-2-2021 14:33:59	Updated	Update Success
23-2-2021 13:57:43	ConfigGenInProgress	ConfigGenInProgress
23-2-2021 13:57:44	ConfigGenDone	Config generation done
1-3-2021 15:21:44	Updated	Update Success
1-3-2021 15:31:48	ConfigGenInProgress	ConfigGenInProgress
1-3-2021 15:31:48	ConfigGenDone	Config generation done

Count: 7

**CLOSE**

The **More actions** () menu at the upper right provides access to the usual list management controls, as described in [Lists](#).

To view the workload intent event log:

#### Procedure

- Step 1.** Open a workload intent as described in [Viewing a workload intent](#).
- Step 2.** Click the **More actions** icon () to open the actions list.
- Step 3.** Select **Event Log** from the drop-down menu.
- Step 4.** Click  at the upper right of the overlay, or the **CLOSE** button, to return to the **Workload Design** view.

## 6.4 Creating and managing QoS profiles

QoS profiles allow you to shape the traffic managed by a workload intent, as described in [QoS profiles](#).

In the Fabric Services System UI, you can:

- Create a QoS profile
- Edit a QoS profile
- Delete a QoS profile

### 6.4.1 Creating a QoS profile

#### Procedure

**Step 1.** If you are not already on the **Profiles** page, do the following:

- Click  to open the main menu.
- Select  **Profiles**.

#### Expected outcome

The **Profiles** page displays.

**Step 2.** From the **Profiles** drop-down, select **QoS**.

**Step 3.** Click **+ CREATE A QOS PROFILE**.

**Step 4.** Enter general information about the QoS profile.

- a. Enter a **Name** for the profile.
- b. Optionally, enter a **Description**.

**Step 5.** Enter **Classifier** data for the QoS profile.

- a. In the **Classifiers** panel, click **+ DSCP POLICY**.

#### Expected outcome

The **Add DSCP Policy** dialog displays.

- b. Select a **Value** and **Forwarding Class** for the policy.
- c. Click **ADD**.

#### Expected outcome

The system adds the DSCP policy to the Classifiers list.

- d. If required, you can delete an individual DSCP policy from the list by clicking on the **More actions** icon (  ) at the right of the policy's row, and selecting **Delete...** from the displayed list.

Repeat this step until the Classifiers list is complete.

**Step 6.** Enter Rewrite Rules for the QoS profile.

Repeat this step until the DSCP Policy list is complete.

- a. In the **Rewrite Rule** panel, click **+ QOS MAPPING**. The **Add QoS Mapping** dialog displays.
- b. Select a **Forwarding Class** and **DSCP** for the mapping entry.

- c. Click **ADD**. The system adds the DSCP policy to the Classifiers list.
- d. If required, you can delete an individual DSCP policy from the list by clicking on the more icon (  ) at the right of the policy's row, and selecting **Delete...** from the displayed list.

Note that the **Workload Reference List** is empty. This list shows all of the workload intents that are currently using this global QoS profile; but because this is a new profile, no workload intents are using it.

**Step 7.** At the lower right of the QoS overlay, click **CREATE**.

#### **Expected outcome**

The system closes the **Create QoS Profile** overlay and returns you to the **Profiles** page with the **QoS** view selected. The profile you just defined is now included in the list of available QoS profiles.

## 6.4.2 Editing a QoS profile

### **About this task**

You can edit a QoS profile at any time.

If the QoS profile you edit is already being used by a workload intent, and that workload intent has already been generated or deployed, then the changes you make do not propagate to the workload intent until you regenerate that workload intent:

- If the workload intent has been generated, but is not yet deployed, you can re-save and regenerate the workload intent without creating a new version. Regenerating the workload intent incorporates the new QoS settings into its configuration.
- If the workload intent has already been deployed, you must create a new candidate version of the workload intent before you can regenerate and redeploy it with the new QoS settings.

### **Procedure**

**Step 1.** If you are not already on the **Profiles** page, do the following:

- a. Click  to open the main menu.
- b. Select  **Profiles**.

**Step 2.** From the **Profiles** drop-down, select QoS.

**Step 3.** Select a QoS profile from the list.

**Step 4.** Click the **More actions** icon (  ) at the right edge of the row, and select **Open** from the drop-down list.

**Step 5.** Update parameters in the QoS profile as described in [Creating a QoS profile](#).

**Step 6.** At the lower right of the **QoS** overlay, click **SAVE**.

### 6.4.3 Deleting a QoS profile

#### About this task

You can delete a QoS profile provided it has not been used in a workload intent. If you try to delete a QoS profile that is in use, the system prevents the action and displays a message indicating why the profile cannot be deleted.

#### Procedure

**Step 1.** If you are not already on the **Profiles** page, do the following:

- a. Click  to open the main menu.
- b. Select  **Profiles**.

**Step 2.** From the **Profiles** drop-down, select **QoS**.

**Step 3.** Select a QoS profile from the list, click the **More actions** icon (  ) at the right edge of the row, and select **Delete...** from the drop-down list.

**Step 4.** In the confirmation dialog, click **OK**.

#### Expected outcome

The system deletes the selected profile and closes the confirmation dialog, returning you to the **Profiles** page with the **QoS** view selected. The profile you just deleted no longer appears in the list.

## 6.5 Creating and managing match groups

Match groups allow you specify a profile for specific types of packets which can then be used to indicate their inclusion or exclusion from workload traffic as described in [ACL profiles](#).

In the Fabric Services System UI, you can:

- Create a match group
- Edit a match group
- Delete a match group

### 6.5.1 Creating a match group

#### Procedure

**Step 1.** If you are not already on the **Profiles** page, do the following:

- a. Click  to open the main menu.
- b. Select  **Profiles**.

**Step 2.** From the **Profiles** drop-down, select **Match Groups**. The system displays the **Match Groups** page, showing a list of previously created match groups.

**Step 3.** To create an IPv4 or IPv6 match group, do one of the following:

- To create an IPv4 match group, go to step [4](#).
- To create an IPv6 match group, go to step [8](#).

**Step 4.** Click **+ CREATE IPV4 MATCH GROUP**.

**Step 5.** Enter general information about the match group:

- a. Enter a **Name** for the match group.
- b. Optionally, enter a **Description**.

**Step 6.** Enter IPv4 Match Entry information for the match group:

- a. In the **IPv4 Match Entry** panel, click **+ADD**.
- b. Enter an IP address in the resulting dialog.  
The IP address must be specified as a prefix; that is, the host section must be all zeros.
- c. Click **ADD**.
- d. Repeat steps [6.a](#) through [6.c](#) until the IPv4 Address list is complete.  
Note that the **ACL Reference List** is empty. This list shows all of the ACL policies that are currently using this IPv4 match group; but because this is a new match group, no profiles are using it.
- e. At the lower right of the Match Group overlay, click **CREATE**.

**Expected outcome**

The system closes the Match Group Creation overlay and returns you to the Profiles page with the Match Group view selected. The match group you just created is now included in the list of available IPv4 match groups.

Repeat this step until the IPv4 Match Entry list is complete.

**Step 7.** Do one of the following:

- To create an IPv6 match group, go to step [8](#).
- If you are finished creating match groups, go to [11](#).

**Step 8.** Click **+ CREATE IPV6 MATCH GROUP**. The **Match Group Creation** overlay displays.

**Step 9.** Enter general information about the match group:

- a. Enter a **Name** for the match group.
- b. Optionally, enter a **Description**.

**Step 10.** Enter IPv6 Match Entry information for the match group:

- a. In the IPv6 Match Entry panel, enter an IPv6 addresses by clicking **+ADD**.
- b. Enter an IP address in the resulting dialog.  
The IP address must be specified as a prefix; that is, the host section must be all zeros.
- c. Click **ADD**.
- d. Repeat steps [10.a](#) through [10.c](#) until the IPv6 Address list is complete.  
Note that the ACL Reference List is empty. This list shows all of the ACL policies that are currently using this IPv6 match group; but because this is a new match group, no profiles are using it.
- e. At the lower right of the Match Group overlay, click **CREATE**.

**Expected outcome**

The system closes the **Match Group Creation** overlay and returns you to the **Profiles** page with the **Match Group** view selected. The match group you just created is now included in the list of available match groups.

**Step 11.** You have completed this procedure.

## 6.5.2 Editing a match group

**About this task**

You can edit a match group at any time.

After you edit a match group, you must update any ACLs that rely on that match group. To aid you in identifying the affected ACLs, these ACLs display a True flag in their "Need update" status. Open and save the ACL.

If the updated ACL profile is being used by a workload intent, and that workload intent has already been generated or deployed, then you must regenerate that workload intent:

- If the workload intent has been generated but is not yet deployed, you can re-save and regenerate the workload intent without creating a new version. Regenerating the workload intent incorporates the new ACL settings into its configuration.
- If the workload intent has already been deployed, you need to create a new candidate version of the workload intent before you can regenerate and redeploy it with the new ACL settings.

**Procedure**

**Step 1.** If you are not already on the **Profiles** page, do the following:

- a. Click  to open the main menu.
- b. Select  **Profiles**.

**Step 2.** From the **Profiles** drop-down list, select **Match Groups**. The system displays the **Match Groups** page, showing a list of previously created match groups.

**Step 3.** Select a match group from the list, click the **More actions** icon (  ) at the right edge of the row, and select **Open** from the drop-down list.

**Step 4.** Update parameters for the match group as described in [Creating a match group](#).

**Step 5.** At the lower right of the **Match Groups** overlay, click **SAVE**.

## 6.5.3 Deleting a match group

**About this task**

You can only delete a match group that is not being used by one or more ACL profiles. For any match group, a list of ACL profiles that are using it are listed in the ACL Reference List when viewing the match group's details.

To delete a match group:

**Procedure**

**Step 1.** If you are not already on the **Profiles** page, do the following:

a. Click  to open the main menu.

b. Select  **Profiles**.

**Step 2.** From the **Profiles** drop-down, select **Match Groups**.

**Step 3.** Select a match group from the list by clicking on the **More actions** icon (  ) at the right edge of the row, and select **Delete...** from the drop-down list.

**Step 4.** In the confirmation dialog, click **OK**.

#### **Expected outcome**

The system deletes the selected match group and closes the confirmation dialog, returning you to the **Profiles** page with the **Match Groups** view selected. The match group you just deleted no longer appears in the list.

## 6.6 Creating and managing ACL profiles

ACL profiles allow you specify whether a workload intent's traffic should include or exclude packets that correspond to one or more match groups, as described in [ACL profiles](#).

The Fabric Services System currently allows you to configure ACL profiles as ingress filters only.

The Fabric Services System does not validate the combinations of fields that constitute an ACL entry. You must configure these fields in accordance with the requirements of the target platform (SR Linux). For this reason, you should be aware of the applicable requirements, limitations, and dependencies of the participating nodes when configuring any ACL profile as part of a workload intent.

For example, the Fabric Services System allows you to configure an IPv4 ACL entry with the following attributes:

- Protocol = icmp
- Destination port = 80

However, this configuration is not permitted by SR Linux because port numbers can only be configured for protocols like UDP/TCP.

If you deploy a workload intent that includes an ACL configuration that is unsupported by the target node, the deployment results in an error. A message from the platform indicates the conflicting configuration.

In the Fabric Services System UI, you can:

- Create an ACL profile
- Edit an ACL profile
- Delete an ACL profile

### 6.6.1 Creating an ACL profile

#### **Prerequisites**

Before you create an ACL profile, ensure that you have created match groups that describe the types of packets you want to accept or reject as part of the ACL; see [Creating a match group](#).

## About this task

When you create an ACL profile, you define a set of packets that should be accepted or rejected by the system.

You define these packet sets by selecting one or more previously created match groups. Each match group already defines one set of possible packets properties; the ACL profile assembles the match groups to create a full profile of the packets deemed acceptable for the current profile.

Later, you can assign ACL profiles to workload intents to represent the packets that are acceptable or unacceptable for the workload intent. The workload intent either accepts or rejects, depending on your selection, all packets that conform to the profiles encompassed by its assigned ACLs.

The total number of entries created for a single IPv4/IPv6 ACL is the product of the following numbers:

- the number of IP address in the source match-groups
- the number of IP addresses in destination match-groups

## Procedure

**Step 1.** If you are not already on the **Profiles** page, do the following:

- a. Click  to open the main menu.
- b. Select  **Profiles**.

**Step 2.** From the **Profiles** drop-down, select **ACL**.

**Step 3.** Click **+ CREATE AN ACL PROFILE**.

**Step 4.** Enter general information about the ACL profile:

- a. Enter a **Name** for the ACL profile.
- b. Optionally, enter a **Description**.

**Step 5.** Select the IPv4 match groups that represent packets acceptable for this ACL:

- a. In the **Match Group Mappings IPv4** panel, click **+ADD**.
- b. Select a **Priority** level.
- c. In the **Accept/Reject** drop-down, select either **Accept** or **Reject**.  
This determines whether the Match Group you are selecting is intended to define acceptable, or unacceptable, packet types.
- d. In the **Source Match Groups** panel, click the box to the left of each pre-existing IPv4 match group you want to associate with the packet source.  
The ACL now represents an instruction to either accept or reject (based on your selection above) any packet whose source IP address conforms to the information in the selected match groups.
- e. In the **Destination Match Groups** panel, click the box to the left of each pre-existing IPv4 match group you want to associate with the packet destination.  
The ACL now represents an instruction to either accept or reject (based on your selection above) any packet whose destination IP address conforms to the information in the selected match groups.
- f. Set IPv4 Match Entry values for the packets to be accepted, or rejected, by this ACL:
  - First Fragment
  - Fragment

- Source Port Operator
- Source Port Value  
The Port Value can either be a number or a text string associated with a predefined port.  
Type a number or select a value from the drop-down list.
- Source Port Range Start
- Source Port Range End
- Destination Port Operator
- Destination Port Value
- Destination Port Range Start
- Destination Port Range End
- ICMP Code (can support multiple values)
- ICMP Type
- Protocol
- TCP Flags

**g.** Click the **ADD** button.

**Step 6.** Select the IPv6 match groups that represent packets acceptable for this ACL:

**a.** In the Match Group Mappings IPv6 panel, click **+ADD**.

**Expected outcome**

The **Match Group Mapping IPv6 Details** overlay displays.

**b.** Select a Priority level.

**c.** In the **Accept/Reject** drop-down, select either **Accept** or **Reject**.

This determines whether the Match Group you are selecting is intended to define acceptable or unacceptable packet types.

**d.** In the **Source Match Groups** panel, click the box to the left of each pre-existing IPv6 match group that you want to associate with the packet source.

The ACL now represents an instruction to either accept or reject (based on your selection above) any packets whose source IP address conforms to the information in the selected match groups.

**e.** In the **Destination Match Groups** panel, click the box to the left of each pre-existing IPv6 match group you want to associate with the packet destination.

The ACL now represents an instruction to either accept or reject (based on your selection above) any packets whose destination IP address conforms to the information in the selected match groups.

**f.** Set IPv6 Match Entry values for the packets to be accepted, or rejected, by this ACL:

- Source Port Operator
- Source Port Value  
The Port Value can either be a number or a text string associated with a predefined port.  
Type a number or select a value from the drop-down list.
- Source Port Range Start

- Source Port Range End
  - Destination Port Operator
  - Destination Port Value
  - Destination Port Range Start
  - Destination Port Range End
  - ICMP Code (can support multiple values)
  - ICMP Type
  - Next Header
  - TCP Flags
- g.** Click the **ADD** button. The system adds the Match Group to the Match Group Mappings IPv6 list.

Note that the Workload Reference List is empty. This list shows all of the workloads that are currently using this ACL profile; but because this is a new ACL profile, no profiles are using it.

**Step 7.** Click the **CREATE** button.

## 6.6.2 Editing an ACL profile

### About this task

You can edit an ACL profile at any time.

If the ACL profile that you edit is already being used by a workload intent, and that workload intent has already been generated or deployed, then the changes you make do not propagate to the workload intent until you regenerate that workload intent.

- If the workload intent has been generated but is not yet deployed, you can re-save and regenerate the workload intent without creating a new version. Regenerating the workload intent incorporates the new ACL settings into its configuration.
- If the workload intent has already been deployed, you must create a new candidate version of the workload intent before you can regenerate and redeploy it with the new ACL settings.

### Procedure

**Step 1.** If you are not already on the **Profiles** page, do the following:

- a.** Click  to open the main menu.
- b.** Select  **Profiles**. The Profiles page displays.

**Step 2.** From the **Profiles** drop-down list, select **ACL**. The system displays the ACL page, showing a list of previously created ACL profiles.

**Step 3.** Select an ACL profile from the list, click the **More actions** icon (  ) at the right edge of the row, and select **Open** from the drop-down list. The **ACL Creation** overlay displays.

**Step 4.** Update parameters for the ACL profile as described in [Creating an ACL profile](#).

**Step 5.** At the lower right of the **ACL** overlay, click **SAVE**.

### Expected outcome

The system saves your changes and closes the **ACL Creation** overlay.

### 6.6.3 Deleting an ACL profile

#### About this task

There are some restrictions in place when deleting an ACL profile to ensure that you do not invalidate any fabric intents that rely on it:

- If an ACL profile has been assigned to a bridged subnet or a sub-interface, or to a routed sub-interface, the system prevents you from deleting the ACL.
- You cannot delete an ACL profile associated with a previous version of a deployed workload intent, even if you are designing a subsequent, undeployed version of that same intent that no longer relies on that ACL profile.
- However, after you deploy a workload intent that no longer relies on an ACL profile, the system allows the deletion of the unassociated ACL profile (provided no other workload intent still relies on it).

To delete an ACL profile:

#### Procedure

**Step 1.** If you are not already on the **Profiles** page, do the following:

- a. Click  to open the main menu.
- b. Select  **Profiles**.

**Step 2.** From the **Profiles** drop-down, select **ACL**.

**Step 3.** Select an ACL profile from the list, click the **More actions** icon (  ) at the right edge of the row, and select **Delete...** from the drop-down list.

**Step 4.** In the confirmation dialog, click **OK**.

#### Expected outcome

The system deletes the selected ACL profile and closes the confirmation dialog, returning you to the **Profiles** page with the **ACL** view selected. The ACL profile you just deleted no longer appears in the list.

## 6.7 Creating a workload VPN intent

A workload intent assigns fabric resources to specific sources of demand, as described in [Elements of a workload intent](#).

#### Prerequisites

Before you create a new workload intent, ensure the following:

- The region that will contain the workload intent has been created; see [Deployment regions](#).
- All fabrics that will participate in the workload intent have been created and successfully deployed; see [Fabric intents](#).
- Any QoS profiles you intend to use with this workload intent have been created; see [Creating and managing QoS profiles](#).
- Any ACL profiles you intend to use with this workload intent have been created; see [Creating and managing ACL profiles](#).

- Any LAGs you intend to act as sub-interfaces for your workload intent have already been created within the system; see [Creating LAGs](#).

### Procedure overview

Creating a workload intent involves the following sub-tasks, each consisting of multiple steps:

1. Create the basic workload intent
2. Add subnets to the workload intent
3. Add sub-interfaces to the workload intent

### Next steps

You are now ready to add subnets as described in [Adding subnets to the workload intent](#).

## 6.7.1 Creating the basic workload intent

### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** From the menu, select  **Workload VPN Intents**.
- Step 3.** Click the **+ CREATE A WORKLOAD VPN INTENT** button to display a set of fabric templates. Templates are displayed in a grid view by default. To switch to the list view, select  in the template selection screen. Click  to return to the grid view.
- Step 4.** Click **VPN Template** and click **CREATE**.
 

**Expected outcome**

The Workload VPN Intents page displays in Workload Design view. The left panel of the page shows basic parameters for you to configure.
- Step 5.** Configure basic parameters:
  - a. Enter a **Workload VPN Intent Name**. This name must be unique among all the workload intents managed by the system.
  - b. Optional: Enter a **Description**.
- Step 6.** Choose a **Fabric Intent Type** (either Real or Digital Sandbox).
- Step 7.** Select one or more fabric intents to participate in the workload intent:
  - a. Click the Edit () icon next to **Fabric Intents**. The system opens a list of fabric Intent, filtered to show only deployed fabrics.
  - b. Check the box at the left edge of the row for each fabric you want to include as part of your workload intent.
  - c. Click the **SELECT INTENTS** button. The system closes the Fabric Intents page and returns you to the Workload VPN Intent creation page.
- Step 8.** Click  to save the latest change to the workload design.

### Expected outcome

The display updates to show the selected fabric intent's topology. The system advances the workload intent's Detailed Status to Created and its Version to 1.0.

## What to do next

You are now ready to proceed to [Adding subnets to the workload intent](#).

### 6.7.2 Adding subnets to the workload intent

#### Procedure

**Step 1.** If you are not continuing directly from the procedure [Creating the basic workload intent](#), first open the Workload VPN Intent view by doing the following:

- a. Click  to open the main menu.
- b. From the menu, select  **Workload VPN Intents**.

**Step 2.** In the **View** drop-down, select **Subnets**.

**Step 3.** Do the following for each bridged subnet you want to add to the workload intent:

- a. Click **+CREATE A SUBNET**. The Workload Subnets page displays.
- b. Enter a Name for the subnet.  
Because the workload name is unique, you can re-use subnet names in different workload intents.
- c. Optionally, enter a **Description** for the subnet.
- d. Select a subnet **Type** of Bridged.
- e. Click **+ADD**.

#### Expected outcome

An Add Entry dialog displays.

- f. Optionally, enter an IP address for the subnet's IP Anycast Gateway.  
This IP address acts as an IRB interface. The subnet can span one, two, or more nodes.
- g. Click **ADD**.
- h. Continue to add gateways until the list is complete.  
The system can currently accommodate up to four gateways.
- i. If your subnet includes IPv4 addresses, set the **IPv4 Learn Unsolicited ARP Enabled** drop-down to either True or False.
- j. If your subnet includes IPv6 addresses, set the **IPv6 Learn Unsolicited ARP Enabled** drop-down to either True or False.
- k. In the **ACL Profile** field, select an existing profile that the system should apply to the current subnet's traffic.  
An ACL profile can only be applied to a bridged subnet for which an IP gateway IP address has been configured.
- l. Set the **Workload Specific ACL Optimization** drop-down to either On or Off.  
Setting this value to On causes the system to verify that the selected ACL profile's source IP address range is contained within the set of Gateway IP addresses you enter for this subnet. Setting this value to Off disables this validation check.
- m. Select an **IP Maximum Transmission Unit (MTU)** value.

n. Click the **CREATE** button.

**Step 4.** Do the following for each routed subnet you want to add to the workload intent:

a. Click **+CREATE A SUBNET**.

b. Enter a **Name** for the subnet.

Because the workload name is unique, you can re-use subnet names in different workload intents.

c. Optionally, enter a **Description** for the subnet.

d. Select a subnet **Type** of Routed.

You do not add an IRB IP address here. Later, you connect the routed subnet to a sub-interface which attaches to a VRF instance.

e. Click the **CREATE** button.

**Step 5.** In the **View** drop-down, select Workload Design.

**Step 6.** Click  to save the latest change to the workload design.

### What to do next

You are now ready to proceed to [Adding sub-interfaces to the workload intent](#).

## 6.7.3 Adding sub-interfaces to the workload intent

### Prerequisites

If you intend to select sub-interfaces by their label, you must have assigned labels to the intended sub-interfaces as described in [Assigning labels to an edge link interface](#).

### About this task

A workload sub-interface consists of an edge-link port or LAG with which you associate ACL and QoS policies. Each sub-interface is associated with a previously created subnet.

The Fabric Services System supports two methods for selecting the edge link port or LAG that constitutes a sub-interface:

- Node and Interface: explicitly select a node and then an interface on that node.
- Port Group Label: assign the Edge-Link label to a set of objects, and then select the label from among those previously created and assigned to underlay interfaces. All interfaces with the specified label are selected.

To add one or more sub-interfaces to the workload intent:

### Procedure

**Step 1.** Do one of the following:

- Open the subnet list and click the **More actions** icon (  ) at the right edge of the row. Select **Create Sub-Interface** from the displayed list.
- Select **Sub-Interfaces** from the Workload VPN intent's View menu and then click **+CREATE A SUB-INTERFACE** in the resulting sub-Interfaces page.

**Step 2.** In the Basic Properties panel, do the following:

- a. Optionally, enter a **Description** for the subnet.

- b. Select a subnet **VLAN ID**.
  - c. Select an **ACL Profile**.
  - d. If this is a sub-interface associated with a routed subnet, set the **Workload Specific ACL Optimization** drop-down to either On or Off.  
Setting this value to On causes the system to verify that the selected ACL profile's source IP address range is contained within the set of Gateway IP addresses you enter for this sub-interface.  
Setting this value to Off disables this validation check.
- Step 3.** Do one of the following:
  - To select sub-interfaces by label, go to step 4.
  - To select sub-interfaces by selecting individual nodes and ports, go to step 5.
- Step 4.** In the Associations panel, do the following:
  - a. In the **Association Type** drop-down, select Port Group Label.
  - b. In the **Port Group Label** field, click  to open the Label Picker dialog.
  - c. From the list of labels, locate the "Edge-Link" label you created previously to identify the edge link ports. Click the  from the left end of the row beside the label.
  - d. Click **SELECT** to close the Label Picker dialog.
  - e. Repeat sub-steps 4.b through 4.d until you have selected all of the intended sub-interfaces.
  - f. Go to step 6.
- Step 5.** In the Association panel, do the following:
  - a. Select an existing Subnet.
  - b. Note that the Association Type is set to Node and Interface. This is a read-only value.
  - c. Select a Node ID associated with a leaf node.  
You must select a leaf node here, because only leaf nodes possess the edge link connections required by the eventual workload.
  - d. Select an Interface Name to identify a specific interface on the selected node.
  - e. Click **+ADD** above the **IP Gateway** field.  
The interface you select here can be a LAG, if the LAG has already been provisioned.
  - f. Enter an IP address and click **ADD**.
  - g. Repeat sub-steps 5.e and 5.f until the IP Gateway list is complete.
- Step 6.** In the QoS panel, assign QoS profiles for the following:
  - QoS Classifier IPv4
  - QoS Rewrite Rules IPv4 (only for a routed subnet)
  - QoS Classifier IPv6
  - QoS Rewrite Rules IPv6 (only for a routed subnet)
- Step 7.** Click the **CREATE** button.
- Step 8.** In the View drop-down, select Workload Design.
- Step 9.** Click  to save the latest change to the workload design.

**Step 10.** Click  **GENERATE WORKLOAD.**

**Expected outcome**

The system generates configuration data for the nodes involved in the workload intent and advances the workload State to Configuration Generated. The workload Version remains 1.0.

## 6.7.4 Configuring BGP

### Prerequisites

Because you create BGP within a workload intent, you must have created a workload intent before you configure BGP.

### About this task

Border Gateway Protocol (BGP) is an inter-AS routing protocol. An AS is a network or a group of routers logically organized and controlled by common network administration. BGP enables routers to exchange network reachability information, including information about other autonomous systems that traffic must traverse to reach other routers in another AS.

When you use BGP as the provider edge (PE) or customer edge (CE) routing protocol, you configure external peering between the provider's AS and the customer network AS.

When you create eBGP links between leaf nodes and customer autonomous systems, the customer autonomous systems may learn of routes through the fabric from different sources. The eBGP links created with the Fabric Services System are configured so that a customer AS prefers the route it learns from its local peer, because that is likely the most efficient path. This is achieved using the BGP Local Preference attribute, which the Fabric Services System sets to a value of 130 for links between peers (while other links generally have a preference value of 100). This behavior is automatic and is not configurable.

### Procedure

**Step 1.** Choose one of the following:

- If you are configuring BGP for a workload intent that has not yet been deployed, open the workload intent in Workload Design view and go to step 2.
- If you are configuring BGP for a workload intent that is already deployed, begin by creating a new candidate version of the existing workload intent as described in [Creating a new version of a workload intent](#).

**Step 2.** From the **View** drop-down, select **Routing**.

**Expected outcome**

The Routing page displays, showing a list of nodes within the workload's fabric or fabrics that are available for BGP configuration.

**Step 3.** Select the row of a node on which to configure BGP.

**Step 4.** Click the **More actions** icon (  ) at the right edge of the row and select **Open BGP** from the displayed action menu.

**Step 5.** Create a BGP group.

- To create a BGP group with some default values, go to step 6.
- To create a BGP group and configure all available values manually, go to step 8

**Step 6.** Create the initial PE-CE BGP Group.

- a. In the Workload BGP panel, enter the following global parameters for the workload PE-CE BGP:

- Router ID
- Peer AS
- Local AS

- b. Click **SAVE**.

**Expected outcome**

The system saves the global parameters, and creates a new BGP Group that appears in the list on the BGP Groups panel. This BGP Group is a read-only collection of the BGP configuration parameters you entered, plus some automatic configuration settings.

This group is a prerequisite for the creation of one or more BGP neighbors.

- Step 7.** Go to step 9.

- Step 8.** Create a PE-CE BGP group:

- a. In the Create BGP Groups panel, click + **CREATE BGP GROUP**.
- b. Enter a **Group Name**.
- c. Click the **BFD** toggle to enable or disable bidirectional forwarding for the group.
- d. Enter a **Connect-Retry** value.
- e. Click the **Override Peer AS** toggle to enable or disable this override. If enabled, enter a Peer AS.
- f. Click the **Override Local AS** toggle to enable or disable this override. If enabled, enter a Local AS.
- g. Use the **IPv4 Unicast** drop-down to enable or disable this feature.
- h. Use the **IPv6 Unicast** drop-down to enable or disable this feature.
- i. Enter a **Minimum-Advertisement-Interval** value.

**Expected outcome**

The system saves the global parameters, and creates a new BGP Group that appears in the list on the BGP Groups panel. This BGP Group is a read-only collection of the BGP configuration parameters you entered.

This group is a prerequisite for the creation of one or more BGP neighbors.

- Step 9.** Create a BGP neighbor:

- a. In the BGP Neighbors panel, click + **CREATE BGP NEIGHBOR**.
- b. Enter the following values:
  - Peer Address
  - Local Address
  - Group Name: a default group name is suggested; retain this value or enter a new one.
- c. Configure the following settings:
  - **Override Peer AS:** This override is disabled by default. To enable this override, click the **Override Peer AS** toggle and enter a Peer AS value.
  - **Override Local AS:** This override is disabled by default. Optionally, click the toggle to enable this override and enter a local AS value.

- **Override IPv4 Unicast:** This override is disabled by default. Optionally, click the toggle to enable this override and enter an IPv4 Unicast value.
- **Override IPv6 Unicast:** This override is disabled by default. Optionally, click the toggle to enable this override and enter an IPv6 Unicast value.

d. Click **CREATE**.

**Expected outcome**

The Create BGP Neighbor overlay closes. The new neighbor appears in the list of BGP neighbors on the Create BGP overlay.

**Step 10.** Repeat step 9 until all required BGP neighbors have been created.

**Step 11.** On the Create BGP overlay, click **SAVE**.

**Step 12.** Update the workload intent with the new BGP information:

- a. On the Workload VPN Intents page, click the **View** drop-down and select Workload Design.
- b. Click **GENERATE WORKLOAD**.

**Expected outcome**

The workload data updates to include the new BGP information.

The system also adds default policy information to the workload configuration. To view the new workload in detail, you can view the configuration code directly by following the procedure [Viewing a workload intent as code](#).

## 6.8 Modifying a workload intent

If you have not yet deployed a workload intent, you can freely modify its design as described in [Editing a workload intent](#).

If you have already deployed a workload intent, you must create a new candidate version of the workload intent before you can modify it. This process is described in [Creating a new version of a workload intent](#). When you have updated the new design, you can deploy it to replace the previously deployed version.

You can also delete a workload intent, as described in [Deleting a workload intent](#).

### 6.8.1 Editing a workload intent

**Procedure**

**Step 1.** Open the saved workload intent as described in [Viewing a workload intent](#).

**Step 2.** Edit the workload intent as required, including any or all of the following:

- Basic parameters
- Participating fabric intents
- Subnets and their attributes
- Sub-interfaces and their attributes

**Step 3.** When you are finished editing parameters, click  to save the new configuration.

**Step 4.** Click  **GENERATE WORKLOAD**.

**Expected outcome**

The system generates updated configuration data for the nodes involved in the workload intent and advances the workload State to Configuration Generated. The workload version increments.

## 6.8.2 Creating a new version of a workload intent

**About this task**

Follow this procedure to create a new candidate version of an already-deployed workload intent.

**Procedure**

**Step 1.** Open the workload intent as described in [Viewing a workload intent](#).

**Step 2.** From the opened workload intent, click the  menu at the upper right of the page.

**Step 3.** Select **Create Candidate Version...** from the drop-down menu.

**Expected outcome**

The system creates a new, editable version of the current workload intent, and the version number in the left panel is incremented to the next available number. You can now edit the workload intent's parameters and save the result without affecting the original version.

**Step 4.** When you are satisfied with the new design, click **Generate** to generate a new workload intent diagram and its associated configuration code.

**Step 5.** At any time before deploying the new version, you can discard it and revert to the previously deployed version by doing the following:

a. Click the **More actions** icon (  ) in the **Workload Design** view.

b. Select **Discard Changes....**

**Expected outcome**

A confirmation dialog displays.

c. Click **OK**.

**Expected outcome**

The system discards the changes and reverts the display to the previous version of the workload intent. If you discarded changes for the initial version of a workload intent that was never deployed, the workload intent is deleted entirely.

## 6.8.3 Updating a group of sub-interfaces by label reference

**About this task**

If you want to update information for a group of sub-interfaces in a workload intent, you can make bulk updates to items assigned the same label. For example, when a series of edge link ports are identified with the same "Edge-Link" key type label, you can update the description field for each of the participating ports in one action, rather than updating the description field for each port individually.

To update a group of sub-interfaces:

**Procedure**

**Step 1.** Open the workload intent in which the "Edge-Link" label is applied to the sub-interfaces.

- Step 2.** In the View drop-down, select **Sub-interfaces**.
- Step 3.** Click  **EDIT BY LABEL**.
- Step 4.** In the dialog, select  to specify the label to edit. The **Label Picker** dialog opens.
- Step 5.** In the Label Picker, do the following:
- Locate the "Edge-Link" label you created to identify the set of edge link ports in question.
  - Click the  from the left end of the row beside the label.
  - Click **SELECT** to close the Label Picker dialog. You return to the Edit Sub-Interfaces By Label dialog.
- Step 6.** Select the VLAN ID to match the ID associated with the labels you intend to edit.
- Step 7.** Click **SELECT** to close the Edit Sub-Interfaces By Label dialog.
- Step 8.** In the Basic Properties header, enter new information in one or both of the following fields:
- Description
  - ACL Profile
- Step 9.** Click **SAVE**.
- Expected outcome**  
The edits are saved as a bulk update to the participating sub-interfaces. The **Edit Sub-Interfaces** overlay closes, returning you to the Sub-Interfaces page.
- Step 10.** Optionally, to confirm updates, do the following from the Sub-Interfaces page:
- For a specific row item that was subject to the bulk edit, click the  options menu at the right end of the row.
  - Select **Open**.  
**Expected outcome**  
The **Edit Sub-Interface** overlay appears. The edits are visible in the Description and ACL Profile fields.
  - Click **CANCEL** to exit without making any further edits.

## 6.9 Deleting a workload intent

### About this task

You can delete a workload intent using the following steps.

### Procedure

- Step 1.** From the main menu, select  **Workload VPN Intents**.
- Step 2.** Select a workload intent from the list. Then, click the  icon at the far-right edge of its row.
- Step 3.** Select **Delete** from the action list and click **OK** in the confirmation dialog.

### Expected outcome

The workload moves to the "delete config enqueued" state. This means that the deletion workload has been added to the region's deployment pipeline, but has not yet taken place.

**Step 4.** Go to the region's deployment pipeline by clicking on the  icon and selecting **Deployment Pipeline...** from the resulting menu.

**Step 5.** From the deployment pipeline, deploy the workload.

#### Expected outcome

If deployment of the workload deletion fails, you must redeploy the underlay intent. The failed attempt is reported as follows:

- The **Deployment Pipeline** screen reports Error and the **Status Reason** field displays the reason for the failure.
- In the **Event Log** screen, a message is reported similar to the following:

```
Message: Delete Workload ww1 Deployment Failed. Underlay Intent needs to be redeployed.
```

- In the **Workload VPN Intents** screen, the workload intent for the deletion (created in step 3) is deleted.
- In the **Fabric intents** view, the **Status** field for the workload intent for the deletion shows Deployed.

## 6.10 Deploying a workload intent

Deploying a workload intent creates a functioning instance of the workload intent as an overlay to your fabric.

The following circumstances can prevent you from deploying a workload intent:

- If any of the fabric intents associated with this workload intent are in the process of being altered (implying a fabric intent alteration which began somewhere in the interval since you created the workload intent), you cannot deploy the workload intent. Any attempt to deploy the workload intent from the deployment pipeline fails, displaying an error message about the fabric underlays having been altered.

In such a case, the system requires you to delete the current candidate version of the fabric intents, thereby reverting to the previous version. You can then proceed with the deployment of your workload intent.



**Note:** You can still deploy the workload intent if the fabric intent alteration is under another workload intent, but uses different nodes.

- If any of the fabric intents associated with this workload intent has undergone a deployed topology change, you cannot deploy the workload intent. Any attempt to deploy the workload intent from the deployment pipeline fails, displaying an error message about the fabric underlays having been altered. In such a case, the system requires you to discard or delete the workload.
- If any of the nodes within the associated fabric intent are unavailable (that is, not in a Ready state), you cannot deploy the workload intent. In such a case, you must correct the node state. When all nodes are back in a Ready state, you can proceed with the deployment of your workload intent.
- If any of the nodes that belong to a workload intent are already under deployment by another workload intent, you must wait until the deployment of the previous workload intent has completed.

Before you can deploy your workload intent you must have saved the workload intent and generated its configuration.

Deploying the workload intent involves two procedures:

1. From the workload intent list or the Design view for an individual workload intent, add the workload intent to the region's deployment pipeline.
2. From the deployment pipeline, select the workload intent and select Deploy from the actions menu.

## 6.10.1 Adding a workload intent to the deployment pipeline

### About this task

To add the workload intent to the region's deployment pipeline, do the following:

#### Procedure

- Step 1.** Open the list of workload intents as described in [Viewing a workload intent](#).
- Step 2.** Select the workload intent you want to deploy from the displayed list.
- Step 3.** Click the **More actions** icon (  ) at the right edge of the workload intent's row, and select **Open** from the drop-down menu.
- Step 4.** Click the Deploy icon (  ).
- Step 5.** Click the **ADD TO PIPELINE** button.

#### Expected outcome

The system adds the workload intent to the deployment queue for the region, and updates the workload intent's Status to "Queued for Deployment."

## 6.10.2 Deploying a workload intent from the deployment pipeline

### About this task

After you add a workload intent to the deployment pipeline, it remains there until you tell the system to proceed with the deployment.

When you are ready to proceed with deployment, do the following:

#### Procedure

- Step 1.** Click the  menu.
- Step 2.** Select  **Workload VPN Intents**.

#### Expected outcome

The Workload Intents page displays, showing a list of workload intents already created. You can sort and filter the list of fabric intents as described in [Lists](#).

- Step 3.** Select the row corresponding to the workload intent you want to deploy.
- Step 4.** Click the  icon at the far-right edge of the row and select **Deployment Pipeline...** from the actions list.
- Step 5.** Find your workload intent in the deployment pipeline list.

**Step 6.** At the right edge of the row, click the **More actions** icon (  ).

**Step 7.** From the resulting actions list, select **Deploy**.

**Expected outcome**

If deployment fails, the failure is reported as follows:

- Queue status: reports Error with detailed status reason
- Fabric intent: reports Deployed and shows a new entry in the Event log showing that the Workload deployment failed
- Workload intent: reports Failed

In the workload intent Design view, the system also highlights them as deployment issues in the status bar by adding a red circle to the fabrics affected by the deployment error, and with entries in the Events log.

## 7 Maintenance intents

A maintenance intent identifies a change you want to deploy to a node.

This release of the Fabric Services System supports two types of maintenance intent:

- Software changes, which upgrade or downgrade the SR Linux software version running on one or more nodes within a single fabric.
- Node replacement, which updates the hardware association for an existing node within a fabric to a new piece of matching hardware with a different serial number. The intent also causes the system to download the necessary configuration files to the new node so that it can fully resume the role of the hardware it replaced.

When a deployed maintenance intent is in progress, whether for a software update or node replacement, the system locks the affected fabric intent. No other operations can be carried out on the locked fabric until the maintenance activity is finished.

Unlike fabric intents and workload intents, a maintenance intent is a one-time operation. You cannot make new versions of a maintenance intent or otherwise maintain it successively over time. Each new maintenance operation requires a new maintenance intent.

### 7.1 Viewing a maintenance intent

#### Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** From the menu, select  **Maintenance Intents**.

**Step 3.** To open a specific maintenance intent from the list, do one of the following:

- Double-click the row for that intent.
- Select a row, click the  icon at the right edge of that row, and select **Open** from the displayed action list.

#### Expected outcome

The system displays the selected maintenance intent in its **Design Intent** view.

### 7.1.1 Elements of the maintenance intent Design view

Figure 31: Maintenance intent view

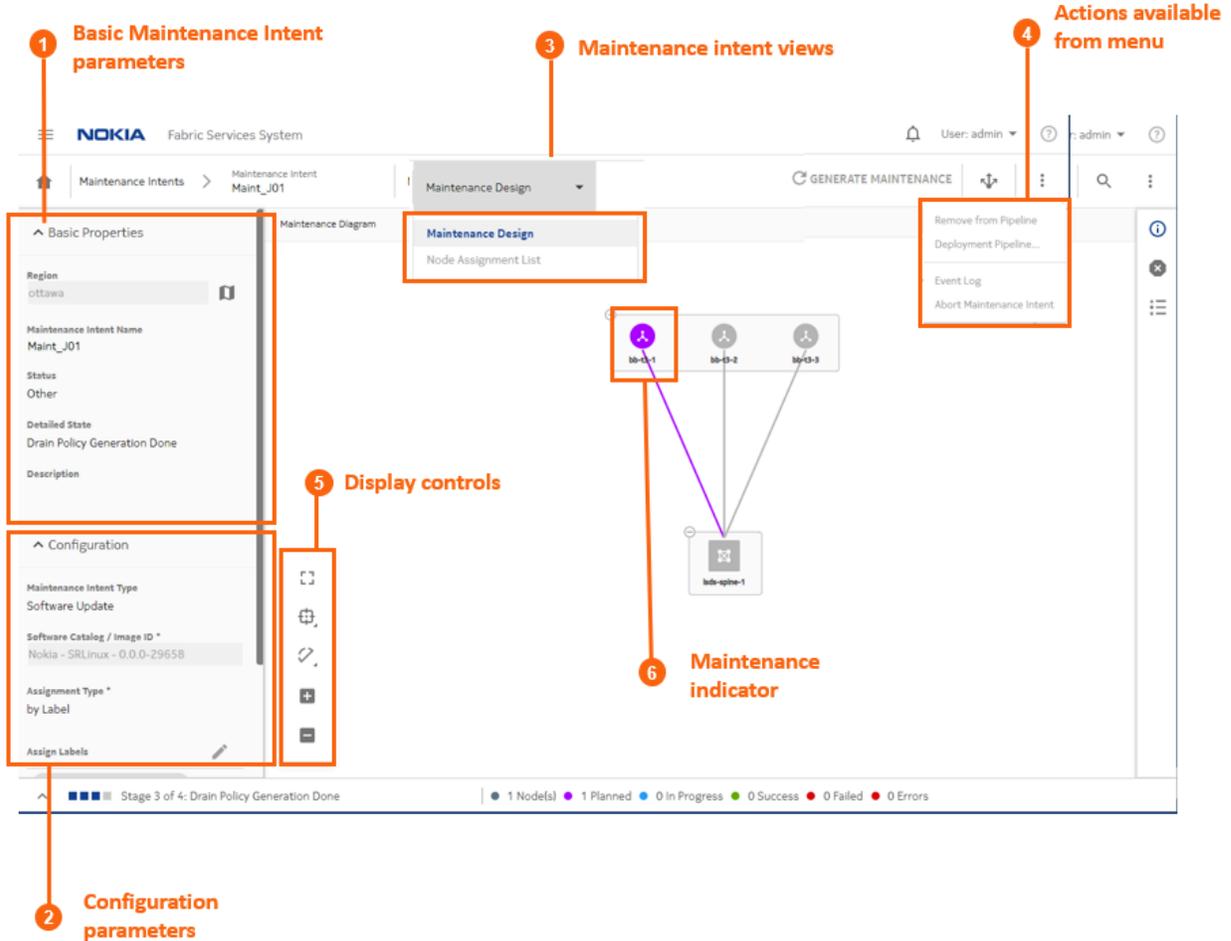


Table 18: Elements of the maintenance intent Design view

	Property	Description
1	Basic Maintenance Intent parameters	This panel shows basic parameters including the intent's name, description, and current state.
2	Configuration parameters	This panel shows the details about the maintenance represented by this intent: <ul style="list-style-type: none"> <li>For software upgrades, this indicates the intended software version and the label that identifies affected nodes.</li> <li>For node replacement, this indicates the new serial number for the node being replaced, and the label that identifies the affected node.</li> </ul>

	Property	Description
3	Maintenance Intent views	The View drop-down provides access to the different views necessary to design a maintenance intent and view a list of the affected nodes.
4	Actions available from menu	From the actions menu you can manage deployment or view the intent's event log.
5	Display controls	These controls allow you to modify the way the system displays the maintenance intent.
6	Maintenance indicator	The purple shading indicates the presence of at least one node that is subject to maintenance.

### 7.1.2 Viewing affected nodes

#### About this task

From an open maintenance intent, you can view a list of all of the nodes that are affected by the current intent.

To view a list of affected nodes:

#### Procedure

- Step 1.** Open a maintenance intent as described in [Viewing a maintenance intent](#).
- Step 2.** From the **Views** drop-down, select Node Assignment List. The system opens a list of the nodes that will be altered by the current maintenance intent.
- Step 3.** To return to the Design view, select **Maintenance Design** from the **Views** drop-down.

### 7.1.3 Viewing the event log

#### About this task

The event log retains a detailed history of all registered events that have occurred with respect to the current maintenance intent. This can be useful when troubleshooting maintenance intent issues, or just to verify that the intent's history matches expectations.

The **More actions** (  ) menu at the upper right of the page provides access to the usual list management controls, as described in [Lists](#).

To view the maintenance intent event log:

#### Procedure

- Step 1.** Open a maintenance intent as described in [Viewing a maintenance intent](#).
- Step 2.** Click the **More actions** icon (  ) to open the actions list.
- Step 3.** Select Event Log from the drop-down menu. The system opens the Event Log overlay.
- Step 4.** Click the  at the upper right of the overlay, or the **CLOSE** button, to return to the maintenance intent Design view.

## 7.2 Creating a maintenance label

### About this task

Before you create a maintenance intent, you must create and apply an identifying label to each of the objects intended to be the target of the intent.

In the current release, these objects are one or more nodes within a single fabric.

Labeling objects is a powerful tool in the Fabric Services System, and is fully described in [Labels](#). In the specific case of labels for a maintenance intent, you create and apply a label with the predefined key "Maintenance," and a value of your choosing.

To create a label you can use to identify a node for maintenance:

### Procedure

**Step 1.** From the main menu, click the  **Label Factory**.

**Step 2.** Click **+ CREATE A LABEL**.

#### Expected outcome

The **Label Details** dialog displays.

**Step 3.** In the **Label Key** drop-down, select Maintenance.

Maintenance is a predefined key that is specifically intended to flag objects as the targets of maintenance intents.

**Step 4.** In the **Value** field, enter text that identifies the maintenance activity.

This value cannot include spaces or special characters.

**Step 5.** Optionally, enter information in the **Comments** field that helps uniquely identify this label.

**Step 6.** Click **CREATE**. The system adds your label to the library of available labels.

## 7.3 Labeling objects for maintenance

### About this task

Before you create a maintenance intent, you must apply an identifying label to each of the objects intended to be the target of the intent.

To label a node for inclusion in a maintenance intent:

### Procedure

**Step 1.** View a node that you intend to be a target of the maintenance operation and open the label assignment dialog using any of the following methods:

- Open the Inventory as described in [Viewing the overall inventory](#) and select the row showing the target node. Then click the **More actions** icon (  ) icon at the right edge of the list and select **Edit Labels** from the resulting menu.
- Open a fabric intent in the **Design and Emulation** view and expand the cluster containing the target node. Right-click the node icon and select **Edit Labels** from the resulting menu.

- Open a fabric intent in the Inventory view, and select the row showing the target node. Then click the **More actions** icon (  ) at the right edge of the list and select **Edit Labels** from the resulting menu.
- Step 2.** In the **Label Picker** dialog, find the maintenance label you created earlier and click the  icon at the left edge of the row.
- Expected outcome**
- The system adds the label to the panel at the top of the **Label Picker** dialog. If other labels were previously applied to the current node, those labels also appear at the top of the **Label Picker** dialog.
- Step 3.** Click **SAVE**. The maintenance label is now associated with the node.
- Step 4.** If your maintenance intent affects multiple nodes, repeat steps 1 through 3 for each of the remaining target nodes.

## 7.4 Creating a maintenance intent

### Prerequisites

Before you create a new maintenance intent, ensure that:

- you have created and applied a maintenance label to the participating nodes as described in [Labeling objects for maintenance](#). The system includes a "Maintenance" key that you can use to identify an object for maintenance. The particular value you associate with that key can be whatever you like.
- If this maintenance intent is for a software update, ensure that the target SR Linux software load is present in the system's software catalog as described in [Software and image catalogs](#).
- If you are using an external DHCP server or file server for SR Linux Zero-touch provisioning (ZTP):
  - for software changes, update the `ztp-config.yml` file with the new software image details.
  - for node replacement, update the `dhcpd.conf` file with the new serial number information.
- If you are performing a software update, ensure that the target node is in a Ready or Version Mismatch state. Attempting the update on a node in any other state results in an error.
- If you are replacing a node, any state is acceptable but some states of the maintenance intent indicating its progress may be skipped depending on the original node state.

### About this task

A maintenance intent indicates a type of change and the nodes this change affects.

To create a maintenance intent:

### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** From the menu, select  **Maintenance Intents**.
- Step 3.** Click the **+ CREATE A MAINTENANCE INTENT** button.
- Step 4.** Configure basic parameters for the maintenance intent:
- a. Select the **Region** to which the maintenance intent belongs

- b. Enter a **Maintenance Intent Name**.
  - c. Optionally, enter a **Description**.
- Step 5.** Do one of the following:
  - If you are creating a maintenance intent for a software update, go to step 6.
  - If you are creating a maintenance intent to replace a node, go to step 7.
- Step 6.** Enter parameters for the software update:
  - a. In the **Maintenance Intent Type** drop-down, select **Software Update**.
  - b. In the **Software Catalog/Image ID** field, select the target SR Linux software load.
  - c. In the **Assign Labels** panel, click the  icon and select the label corresponding to this maintenance action from the displayed list.  
In the current release, you can select only one label as the target for the maintenance action.
  - d. Go to step 9.
- Step 7.** Enter parameters for the node replacement:
  - a. In the **Maintenance Intent Type** drop-down, select **Replace Node**.
  - b. In the **Assign Labels** panel, click the  icon and select the label corresponding to this maintenance action from the displayed list.  
In the current release, you can select only one label as the target for the maintenance action.  
The system updates the **Nodes List** panel to show the node bearing the selected label.
- Step 8.** Enter the replacement node's serial number in the **New Serial Number** field.
- Step 9.** Click the Save icon () to save the maintenance intent.
- Step 10.** Click  **GENERATE MAINTENANCE**.
- Step 11.** Optionally, review the generated configuration code to confirm that it is satisfactory by doing the following:
  - a. Right-click the node that is the target of the maintenance intent and select **Inspect Configuration** from the pop-up menu.
  - b. Review the code and verify that it is satisfactory.
  - c. Close the **Configuration Inspector** by clicking the X at the upper right of the overlay.

### What to do next

You are now ready to deploy the maintenance intent.

## 7.5 Duplicating a maintenance intent

### About this task

You can duplicate an existing maintenance intent, which can save you several steps if you need to apply the same or a similar maintenance action again (to the same or different target nodes).

To duplicate a maintenance intent:

## Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** From the menu, select  **Maintenance Intents**.

**Step 3.** Click a maintenance intent in the list to select it, and then do one of the following:

- Double-click the row for that intent.
- Select a row, click the  icon at the right edge of that row, and select **Open** from the displayed action list.
- Click the  menu at the upper right of the page and select **Abort Maintenance Intent**.

Any further action on the maintenance intent is prevented, and the target node is restored to its configured state before the maintenance intent was applied.

**Step 4.** Click the **More actions** icon () at the right edge of that row, and select **Duplicate Intent** from the displayed action list.

### Expected outcome

The Maintenance Intents page displays, with the same data as the original intent except for **Maintenance Intent Name**, which is blank.

**Step 5.** Enter a name for the new, duplicate maintenance intent.

**Step 6.** Update other values in the maintenance intent if required.

**Step 7.** Continue to save, generate, and deploy the maintenance intent as described in [Creating a maintenance intent](#).

## 7.6 Deploying a maintenance intent

Deploying a maintenance intent causes the system to deploy configuration files to the affected nodes, which ultimately results in the fabric changes embodied in the intent.

Before you can deploy your maintenance intent you must have saved the maintenance intent and generated its configuration.

Deploying the maintenance intent involves two procedures:

- Adding the maintenance intent to the deployment pipeline causes the intent to take its place in the list of planned deployments of fabric, workload, and maintenance intents for the current region. Adding a maintenance intent to the pipeline triggers certain preparatory actions that the system performs behind the scenes. These actions must finish before the system allows you to proceed with deployment.

For example if your maintenance intent is replacing a node, adding the intent to the deployment pipeline causes the system to stop the DHCP client application on the existing node to releases its current IP address. Only after this is complete can you deploy the maintenance intent from the deployment pipeline.

- From the deployment pipeline you must then manually deploy the intent. This signals to the system that it can proceed with the deployment of configuration files to all participating nodes. This deployment may not occur immediately; the system resolves each active deployment in sequence, waiting until one is complete before proceeding to the next.

## 7.6.1 Adding a maintenance intent to the deployment pipeline

### About this task

To add the maintenance intent to the region's deployment pipeline, do the following:

### Procedure

- Step 1.** Open the list of maintenance intents as described in [Viewing a maintenance intent](#).
- Step 2.** Select the maintenance intent you want to deploy from the displayed list.
- Step 3.** Click the **More actions** icon (  ) at the right edge of the maintenance intent's row, and select **Open** from the drop-down menu.
- Step 4.** Click the **Add Intent to Deployment Pipeline** icon (  ).
- Step 5.** Click the **ADD TO PIPELINE** button.

## 7.6.2 Progress of a deployed maintenance intent

When a deployed maintenance intent is in progress, whether for a software update or node replacement, the system locks the affected fabric maintenance. No other operations can be carried out on the locked fabric until the maintenance activity is finished.

The deployed maintenance intent progresses through a series of states, depending on the type of maintenance being performed.

For a software update:

- **Deploy drain policies (1/3):** the system diverts traffic away from the node in anticipation of the software update or hardware replacement.  
If the maintenance intent fails at this stage, you must discard the fabric intent changes manually. The details of any failure can be found in the Maintenance Error/Event log.  
If this step is successful, the system automatically proceeds to step 2/3.  
**ZTP (2/3):** After the maintenance intent's configuration change has been made upon the node, the system rediscovers the updated node using Zero-touch Provisioning (ZTP). The end result of this state is the post-maintenance node returned to a ready state. If this step is successful, the system automatically proceeds to step 3/3.  
**Deploy undrain policies (3/3):** the traffic previously diverted away from the node is restored.
- If the maintenance intent fails at this stage, you must re-deploy the affected fabric intent manually. The details of any failure can be found in the Maintenance Error/Event log.

For a node replacement:

- **Deploy drain policies (1/3):** the system diverts traffic away from the node in anticipation of the software update or hardware replacement.  
If this step is successful, the system automatically proceeds to step 2/3.  
**Releasing DHCP IP (2/3):** The node's previous IP address is released from DHCP. If this step is successful, the system automatically proceeds to step 3/3.
- **Deploy undrain policies (3/3):** the traffic previously diverted away from the node is restored.  
If the maintenance intent fails at this stage, you will need to re-deploy the affected fabric intent manually. The details of any failure can be found in the Maintenance Error/Event log.

If any issues arise during deployment, they appear as alerts in the alert panel at the right side of the page. In the maintenance intent Design view, the system also highlights them as deployment issues in the status bar, by adding a red circle to the fabrics affected by the deployment error, and with entries in the Events Log.

### 7.6.3 Deploying a maintenance intent from the deployment pipeline

#### About this task

The maintenance intent remains in the deployment pipeline until you tell the system to proceed with the deployment.

Adding a maintenance intent to the pipeline triggers preparatory actions that the system performs behind the scenes. These actions must finish before the system allows you to proceed with deployment.

When you are ready to proceed with deployment, do the following:

#### Procedure

- Step 1.** In the Fabric Services System UI, click the  menu.
- Step 2.** Select  **Maintenance Intents**.
- Step 3.** Select the row corresponding to the maintenance intent you want to deploy.
- Step 4.** Click the  icon at the far-right edge of the row and select **Deployment Pipeline...** from the actions list.
- Step 5.** Find your maintenance intent in the deployment pipeline list.
- Step 6.** At the right edge of the row click the **More actions** icon (  ).
- Step 7.** From the resulting actions list, select **Deploy**.

## 7.7 Removing a maintenance intent from the deployment pipeline

#### About this task

To remove the maintenance intent from the region's deployment pipeline before you have deployed it, do the following:

#### Procedure

- Step 1.** Open the list of maintenance intents as described in [Viewing a maintenance intent](#).
- Step 2.** Select the maintenance intent you want to remove from the deployment pipeline.
- Step 3.** Click the **More actions** icon (  ) at the right edge of the maintenance intent's row, and select **Open** from the drop-down menu.
- Step 4.** Click the **More actions** icon (  ) at the upper right of the **Maintenance Design** page and select **Remove from Pipeline**.

## 7.8 Aborting a maintenance intent

### About this task

If a maintenance intent fails to complete properly, you can abort the deployed maintenance intent. When aborted, any further action on the maintenance intent is prevented, and the target node is restored to its configured state before the maintenance intent was applied.

You can only abort a maintenance intent after a ten-minute interval has passed since the maintenance intent adopted either the "ZTP" or "Waiting for Dial Home" state.

To abort a maintenance intent:

### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** From the menu, select  **Maintenance Intents**.
- Step 3.** To open a specific maintenance intent from the list, do one of the following:
  - a.** Double-click the row for that intent.
  - b.** Select a row, click the  icon at the right edge of that row, and select **Open** from the displayed action list.
  - c.** Click the  menu at the upper right of the page and select **Abort Maintenance Intent**.
  - d.** Click OK.

### Expected outcome

If fewer than 10 minutes have passed, a dialog displays rejecting the Abort instruction. Otherwise, the system cancels the maintenance activity.

## 8 Labels

Labels are tags that help you group and organize fabric items according to specific criteria. You can assign the labels to various fabric items, such as fabric intents, workload VPN intents, maintenance intents, or to objects within each of these intents. For instance, within a fabric intent, you can label specific groupings of nodes. Within a workload intent, you can label specific groupings of subnets and sub-interfaces.

Labels are a powerful tool for identifying a group of objects that can then be subject to collective actions. For example, as part of a maintenance intent, you can create a label and apply it to a group of items that will be subject to a software upgrade. Then, you can apply the upgrade to all of the nodes with that label as a single action, instead of upgrading each node individually. See [Labeling objects for maintenance](#) for a detailed procedure.

Labels can also be used for informational purposes. You can label specific fabric elements based on criteria such as their physical location. For instance, you can create a label to identify nodes located in a specific data center (DC), then apply this label to nodes physically located in that DC.

### 8.1 Label types

You can associate two types of labels: Nokia pre-defined labels and user-configured custom labels. Fabric Services System comes with a series of predefined labels that you can associate with fabric items. If these labels do not meet your specific tagging criteria, you can also create new labels in the Label Factory.

#### 8.1.1 Nokia pre-defined labels

Nokia pre-defined labels are a basic set of grouping labels for common use. Some pre-defined grouping labels have a built-in priority to resolve when conflicting configuration profiles are applied to it.

Pre-defined fabric intent labels are key and value pairings that can be applied to fabric intents and workload intents. Nokia has defined these labels for use in common tagging scenarios. [Table 19: Pre-defined labels](#) lists the Nokia pre-defined labels.

*Table 19: Pre-defined labels*

Label key	Value	Description
Node-type	T1_LEAF	Identifies all T1 nodes in a fabric. It is dynamic, so new T1 nodes can join this group.
	T2_SPINE	Identifies all T2 nodes in a fabric. It is dynamic, so new T2 can join this group.
	T3	Identifies all T3 nodes in a fabric. It is dynamic, so new T3 nodes can join this group.
	T4	Identifies all T4 nodes in a fabric. It is dynamic, so new T4 nodes can join this group.

Label key	Value	Description
	T5	Identifies all T5 nodes in a fabric. It is dynamic, so new T5 nodes can join this group.
Link-type	T1_ISL_T2	Identifies inter-switch links between T1 and T2 nodes. This label is fabric-scoped. It is dynamic so new ISLs can join this group.
	T2_ISL_T3	Identifies inter-switch links between T2 and T3 nodes. This label is fabric-scoped. It is dynamic so new ISLs can join this group.
	T3_ISL_T4	Identifies inter-switch links between T3 and T4 nodes. This label is fabric-scoped. It is dynamic so new ISLs can join this group.
	T4_ISL_T5	Identifies inter-switch links between T4 and T5 nodes. This label is fabric-scoped. It is dynamic so new ISLs can join this group.
	EDGE LINK	Identifies all edge links. This label is fabric-scoped. It is dynamic so new edge links can join this group.

### 8.1.2 User-configured custom labels

User-configured custom labels allow you to group entities on which configuration profiles can be applied. You can define labels that broadly identify fabric elements, such as by geographical region. You can also define labels to identify highly specific elements, such as a physical location within a DC.

Every custom label has two parts:

- a key, which identifies the type of information the label contains
- a value, which is a specific value appropriate to the label's type

When a label is created, the key and value must be assigned, along with a description to describe the purpose of the label. You may want to define a label to identify physical fabric elements found in a specific location. When creating the key- and-value pairing, you can identify the general location as the label key, and a specific region indicator as the value.

Fabric Services System provides some pre-defined label keys that you can use when configuring a custom label. These pre-defined label keys identify common elements of network fabrics that you may want to label and can help to guide you when you create new labels. Consider using any of the label keys as a starting point when you configure a custom label. [Table 20: Pre-defined label keys](#) lists the pre-defined label keys. Nokia recommends using these label keys to ensure consistent naming conventions across fabric elements in your network.

*Table 20: Pre-defined label keys*

Label key	Description	Examples
ACL-GRP	Use to identify a group of switches, downlink ports, or a group of sub-interfaces used to apply ACL policies on the group.	ACL-GRP-1, ACL-GRP-WEB, ACL-GRP-APP, ACL-GRP-DB

Label key	Description	Examples
AZ	Use to identify different availability zones (AZs) in a data center. By dividing the network infrastructure inside a DC into multiple fault boundaries, application teams can place compute across AZs and achieve higher availability to their distributed applications.	AZ-1, AZ-2
Building	Use to identify individual buildings that make up a data center. If the data center exists in a single building, use the DC label type instead.	Building-A, Building-B
DC	Use to identify specific data centers.	DC-1, DC-Europe, DC-Atlanta
Edge-Link	Use to identify downlink ports to be used in a workload design.	Edge-1, Edge-A
Link-type	Use to identify a group of links of a specified type.	ISL-1, ISL-A
Maintenance	Use to identify nodes to be part of a collective maintenance action; either upgrade or node replacement.	Maintenance-upgrade1, Maintenance-replacement1
Node-type	Use to identify a group of nodes of a specified type.	Node-1, Node-A
QOS-GRP	Use to identify a group of switches, downlink ports, or a group of sub-interfaces used to apply QoS policies on the group.	QOS-GRP-1, QOS-GRP-WORKLOAD1, QOS-GRP-MGMT, QOS-GRP-STORAGE
Region	Use to identify a group of resources on a per customer region. A region is a geographical area location, such as a city.	Region-A, Region-EAST, Region-Dallas
Room	Use to identify specific rooms within a data center.	Room-A, Room-B
Tenant	Use to identify a group of switches, downlink ports, or a group of sub-interfaces that describe a tenant in a DC.	Tenant-1, Tenant-ABC

For example, you can create a label to identify fabric elements located in particular Availability Zones (AZs). When entering the key-and-value pairing, you could enter *AZ* as the key, and *Zone1* as the value. The label can be assigned to all fabric elements in this specific zone. You can then create additional availability zone labels with the same *AZ* key, but define additional values to specify additional zones, such as *Zone2*, *Zone3*, and so on.

## 8.2 The Label Factory

The Label Factory shows an inventory of all the Nokia pre-defined labels as well as any user-configured custom labels. From the Label Factory, you create new labels and manage existing labels.

You can also query the Label Assignment list to find a specific subset of fabric items that are assigned a specific label. A label can have hundreds or even thousands of assignments to various fabric items throughout the system.

## 8.3 Viewing available labels

### About this task

Follow this procedure to view a list of all available labels in the Label Factory. The Label Factory opens to the Available Labels category by default. If you navigate away from this list, you can select **Available Labels** from the drop-down list to return to the Available Labels category.

### Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** Select  **Label Factory**.

### Expected outcome

The Label Factory page opens, showing a list of both Nokia pre-defined and user-configured custom labels.

## 8.4 Creating a label

### About this task

You can create a custom label and assign it to one or more fabric elements. Follow this procedure to create a label in the Label Factory.

### Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** In the main menu, select  **Label Factory**.

**Step 3.** Click **+ CREATE A LABEL**.

**Step 4.** In the Label Key drop-down, select one of the pre-defined keys.  
The label keys are described in [Table 20: Pre-defined label keys](#) .

**Step 5.** In the Value field, enter text that identifies the purpose of the label.  
Enter a value that you can easily recognize later. This value cannot include spaces or special characters.

**Step 6.** Optionally, enter information in Comments that helps uniquely identify this label.

**Step 7.** Click **CREATE**.

## 8.5 User-configured label manipulation

After you or another user creates a label, you can edit the description of that label, or delete the label from the available labels entirely.



**Note:** You cannot delete a Nokia pre-defined label, nor edit such labels other than to modify the Comments field.

## 8.5.1 Editing an existing label

### About this task

You can edit an existing label to modify the comments that describe its purpose. Follow this procedure to edit an existing label in the Available Labels list.

### Procedure

- Step 1.** Open the list of available labels as described in [Viewing available labels](#).
- Step 2.** For a specific row item, click the  options menu at the right end of the row.
- Step 3.** Click **Edit Label**.
- Step 4.** Edit the value in the **Comments** field.
- Step 5.** Click **SAVE**.

## 8.5.2 Deleting an existing label

### About this task

Only user-configured labels can be deleted. Before deleting a label, you must remove any associations the label has with fabric objects. For fabric intent labels, see the following topics: [Removing a label assigned to a fabric intent](#) or [Removing a label assigned to a node in a fabric intent](#). For workload intent labels, see the following topics: [Removing a label assigned to a workload intent](#) or [Removing a label assigned to a specific sub-interface](#).

Follow this procedure to delete a label.

### Procedure

- Step 1.** Open the list of available labels as described in [Viewing available labels](#).
- Step 2.** For a specific row item, click the  options menu at the right end of the row.
- Step 3.** Select **Delete...** from the action list.
- Step 4.** Click **OK**.

## 8.6 Label assignments to fabric intent elements

You can assign labels from the Label Factory to various fabric elements. You can assign Nokia pre-defined labels, user-configured labels, or a combination of both. Labels can be assigned to a specific fabric intents, or to specific elements within each intent.

Multiple labels can be associated with a single fabric element. They can also be applied to real or planned nodes in a fabric.

After a label has been assigned to a fabric element, you can also remove it.

## 8.6.1 Assigning labels to a fabric intent

### About this task

When you assign a label to an intent, the individual elements that make up the intent do not inherit the label. These elements can each be assigned their own labels separately.

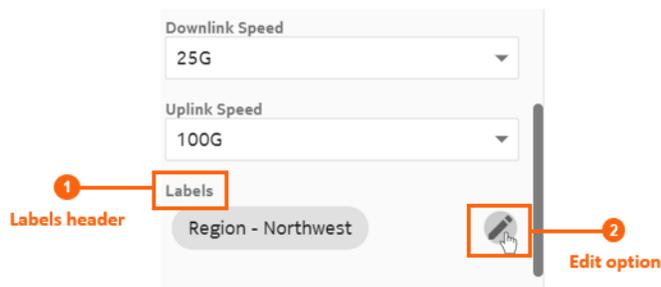
Follow this procedure to assign labels found in the Label Factory to a specific fabric intent.

### Procedure

**Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).

**Step 2.** On the High Level Intent panel, click  from under the Labels header.

Figure 32: Assigning a label to a fabric intent



**Step 3.** From the list of labels, choose one or more labels to assign to the fabric intent. Click **+** from the left end of the row for each label you want to add.

**Step 4.** Click **SAVE**.

## 8.6.2 Assigning a label to a specific node in a fabric intent

### About this task

Follow this procedure to assign labels found in the Label Factory to specific nodes in a fabric intent.

### Procedure

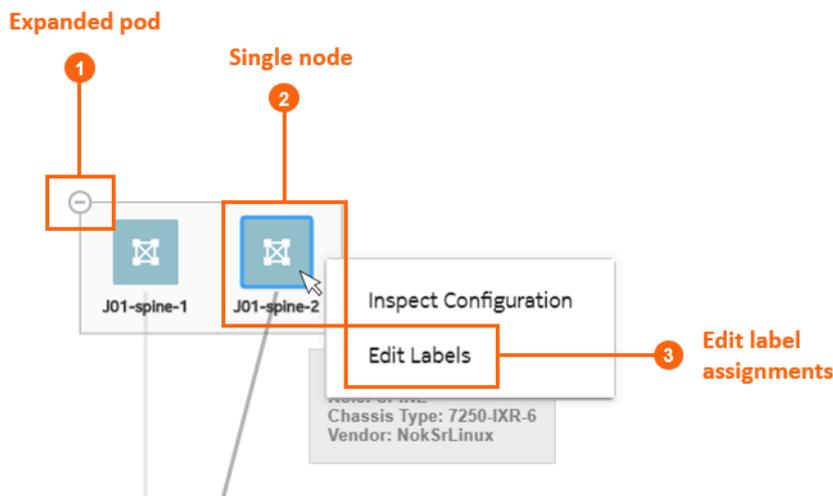
**Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).

**Step 2.** Right-click a node in the cable diagram.

You may have to expand a cluster to find a specific node. See [Groups](#).

**Step 3.** Select **Edit Labels**.

Figure 33: Assigning a label to a specific node



**Step 4.** From the list of labels, choose one or more labels to assign to the fabric intent. Click the **+** from the left end of the row for each label you want to add.

**Step 5.** Click **SAVE**.

### 8.6.3 Assigning labels to a fabric link

#### About this task

Follow this procedure to assign labels from the Label Factory to specific fabric links within the intent.

#### Procedure

**Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).

**Step 2.** Select **Fabric Links** from the drop-down list.

**Step 3.** For a specific fabric link, click the **⋮** options menu at the right end of the row.

**Step 4.** Select **Edit Labels**.

**Step 5.** From the list of labels, choose one or more labels to assign to the fabric link. Click **+** from the left end of the row for each label you want to add.

**Step 6.** Click **SAVE**.

### 8.6.4 Assigning labels to an edge link interface

#### About this task

You can assign labels from the Label Factory to specific edge link interfaces or LAGs within a fabric intent. Edge link interfaces must be labeled using the "Edge-Link" label key. You can create specific labels for edge link interfaces, apply the labels, then reference these labeled items when creating a workload intent sub-interface via the label assignment; see [Adding sub-interfaces to the workload intent](#).

To assign a label to an edge link interface, you must first deploy the fabric intent.

Follow this procedure to assign a label to an edge link interface:

### Procedure

- Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).
- Step 2.** Select **Edge-Links** from the drop-down list.
- Step 3.** For a specific item in the list, click the  options menu at the right end of the row.
- Step 4.** Select **Edit Labels**.
- Step 5.** From the list of labels, choose one or more labels to assign to the edge link port. Click  from the left end of the row for each label you want to add.  
Only labels with the "Edge-Link" label key are displayed.
- Step 6.** Click **SAVE**.

## 8.6.5 Removing a label assigned to a fabric intent

### About this task

Follow this procedure to remove labels assigned to a fabric intent.

### Procedure

- Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).
- Step 2.** On the High Level Intent panel, click  from the Labels section.

### Expected outcome

The Label Picker dialog opens. The currently assigned labels are displayed in the top of the dialog.

- Step 3.** From the assigned labels displayed, click the  next to the label you want to remove from the fabric intent.  
You can do this for one or more labels.
- Step 4.** Click **SAVE**.

## 8.6.6 Removing a label assigned to a node in a fabric intent

### About this task

Follow this procedure to remove labels assigned to node in a fabric intent.

### Procedure

- Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).
- Step 2.** Right-click a node in the cable diagram.
- Step 3.** Select **Edit Labels**.

**Expected outcome**

The Label Picker dialog opens. The currently assigned labels are displayed in the top of the dialog.

**Step 4.** From the assigned labels displayed, click the **X** next to the label you want to remove from the fabric intent.

You can do this for one or more labels.

**Step 5.** Click **SAVE**.

## 8.6.7 Removing a label assigned to a fabric link

**About this task**

Follow this procedure to remove labels assigned to a fabric link.

**Procedure**

**Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).

**Step 2.** Select **Fabric Links** from the drop-down list.

**Step 3.** For a specific fabric link, click the **⋮** options menu at the right end of the row.

**Step 4.** Select **Edit Labels**.

**Expected outcome**

The Label Picker dialog opens. The currently assigned labels are displayed in the top of the dialog.

**Step 5.** From the assigned labels displayed, click the **X** next to one or more labels you want to remove from the fabric link.

**Step 6.** Click **SAVE**.

## 8.6.8 Removing a label assigned to an edge link interface

**About this task**

Follow this procedure to remove labels assigned to an edge link interface.

**Procedure**

**Step 1.** Open a fabric intent as described in [Viewing a fabric intent](#).

**Step 2.** Select **Edge-Links** from the drop-down list.

**Step 3.** For a specific item in the list, click the **⋮** options menu at the right end of the row.

**Step 4.** Select **Edit Labels**.

**Expected outcome**

The Label Picker dialog opens. The currently assigned labels are displayed in the top of the dialog.

**Step 5.** From the assigned labels displayed, click the **X** next to one or more labels you want to remove from the edge link port.

**Step 6.** Click **SAVE**.

## 8.7 Label assignments to workload intent elements

Labels from the Label Factory can be assigned to workload intents, and to the elements contained within workload intents. You can assign Nokia pre-defined labels, user-configured labels, or a combination of both. Labels can be assigned to a specific workload intent, or to specific subnets or sub-interfaces within an intent.

Multiple labels can be associated with a single workload intent or element. You can also remove labels from workload intent elements after they have been assigned.

### 8.7.1 Assigning labels to a workload intent

#### About this task

You can assign labels from the Label Factory to specific workload intents. When you assign a label to an intent, the individual elements that make up the intent do not inherit the label. These elements can each be assigned their own labels separately.

Follow this procedure to assign a label to a workload intent.

#### Procedure

- Step 1.** Open a workload intent as described in [Viewing a workload intent](#).
- Step 2.** In the Workload Design view, on the left panel, click  from under the Labels header.
- Step 3.** From the list of labels, choose one or more labels to assign to the workload intent. Click the **+** from the left end of the row for each label you want to add.
- Step 4.** Click **SAVE**.

### 8.7.2 Assigning a label to a specific sub-interface

#### About this task

You can assign labels found in the Label Factory to specific sub-interfaces within a workload intent. You must have a sub-interface created before beginning this procedure.

Follow this procedure to assign a label to the specific sub-interface.

#### Procedure

- Step 1.** Open a workload intent as described in [Viewing a workload intent](#).
- Step 2.** In the workload intents drop-down list, switch to the **Sub-Interfaces** view.
- Step 3.** For a specific row item, click the  options menu at the right end of the row.
- Step 4.** Select **Edit Labels**.
- Step 5.** From the list of labels, choose one or more labels to assign to the element. Click the **+** from the left end of the row for each label you want to add.
- Step 6.** Click **SAVE**.

### 8.7.3 Removing a label assigned to a workload intent

#### About this task

Follow this procedure to remove labels assigned to a workload intent.

#### Procedure

**Step 1.** Open a workload intent as described in [Viewing a workload intent](#).

**Step 2.** In the Workload Design view, on the left panel, click  from under the Labels header.

#### Expected outcome

The Label Picker dialog opens. The currently assigned labels are displayed in the top of the dialog.

**Step 3.** From the assigned labels displayed, click the  next to the label you want to remove from the workload intent.

You can do this for one or more labels.

**Step 4.** Click **SAVE**.

### 8.7.4 Removing a label assigned to a specific sub-interface

#### About this task

Follow this procedure to remove labels assigned to a specific sub-interface within a workload intent.

#### Procedure

**Step 1.** Open a workload intent as described in [Viewing a workload intent](#).

**Step 2.** In the workload intents drop-down list, switch to the **Sub-Interfaces** view.

**Step 3.** For a specific row item, click the  options menu at the right end of the row.

**Step 4.** Select **Edit Labels**.

#### Expected outcome

The Label Picker dialog opens. The currently assigned labels are displayed in the top of the dialog.

**Step 5.** From the assigned labels displayed, click the  next to the label you want to remove.

You can do this for one or more labels.

**Step 6.** Click **SAVE**.

## 8.8 Label assignment management

In addition to creating labels, you can use the Label Assignments view to understand how the inventory of labels is being used throughout the system. A label can have hundreds or even thousands of assignments, so the Label Assignments page can show you the fabric elements, either real or virtual, that a specific label is assigned to.

## 8.8.1 Viewing the assignments of a specific label

### About this task

You can view a list of resources that each label is assigned to in the Label Factory. Follow this procedure to view the assignments of a specific label.

### Procedure

- Step 1.** Open the available labels as described in [Viewing available labels](#).
- Step 2.** For a specific row item, click the  options menu at the right end of the row.
- Step 3.** Select **Show Assignment**.

### Expected outcome

The Label Assignment page opens to show a list of assignments for the selected label under the Object Name heading.

## 8.8.2 Querying the label assignment list

### About this task

After opening the label assignments of a specific label, you can refine the assignment list to find a specific subset of fabric items with a specific label or series of labels associated. Follow this procedure to query the Label Assignment list.

### Procedure

- Step 1.** Open the Label Factory as described in [Viewing available labels](#).
- Step 2.** Select the **Label Assignment** category filter.
- Step 3.** Enter a search string in the search query field as described in [Creating label assignment list string queries](#).
- Step 4.** Click  **REFRESH LIST**.

### Expected outcome

The list is refined to display only fabric elements that match the query.

- Step 5.** Optionally, add additional details in the search query field to further refine the list of fabric elements.
- Step 6.** Optionally, click  **REFRESH LIST** to display the refined fabric elements list.

### What to do next

You can continue to refine the Label Assignment list by adding more values to the search query.

Click **CLEAR ALL** to clear the queries.

You can view more information about specific fabric elements that have been assigned the label by clicking an element in the list, then expanding the information panel on the right side of the page by clicking . The information panel shows more details about the currently selected fabric element.

### 8.8.3 Creating label assignment list string queries

The Label Assignment page shows a list of all assignments for a specific label. A label can have hundreds or even thousands of assignments.

You can refine the Label Assignment list to find a specific subset of fabric items that are assigned a specific label. Use the search query field to refine the list.

In the search query field, enter an expression to include or exclude particular label values. Operators you can use in this expression include the following (followed by an example for each):

- OR, AND, and NOT

```
"<LabelName1>" = "<LabelValue1>" or "<LabelName2>" = "<LabelValue2>"
```

- IN and NOT IN

```
"<LabelName>" in ( "<LabelValue1>" or "<LabelValue2>" or "<LabelValue3>" )
```

- EXISTS and NOT EXISTS

```
exists ("<LabelName1>" or "<LabelName2>" or "<LabelName3>")
```

Entering a search query allows you to filter a series of fabric items that have specific multiple labels assigned, while excluding fabric items that also have specific other labels assigned.

For example, if you want to see a list of fabric elements assigned the Label\_A label, but exclude any fabric items that are also assigned the Label\_B label, you can enter the following search query:

```
INCLUDE (Label = "Label_A") EXCLUDE (Label = "Label_B")
```

Based on the search query, the Label Assignment list displays only fabric elements assigned Label\_A. The fabric elements in the list may also be assigned multiple other labels. You can continue to refine the search by adding additional criteria to the query.

You can filter the list of available labels as described in [Lists](#).

## 9 Inventories

The Inventory page shows a complete list of all fabric elements across all fabrics in Fabric Services System. When you create a new fabric intent, you specify the number of nodes it should contain (for example, flexible leaf/spine nodes or backbone nodes). These nodes are represented in the inventory. You can also add additional nodes to an existing fabric intent, which also appear in the inventory. The types of nodes in the inventory list include:

<b>planned nodes</b>	Created when you define a fabric intent, but do not yet have associated real-world hardware.
<b>real-world hardware</b>	Physical nodes that are present in the network, but are not associated with planned nodes.
<b>associated nodes</b>	Where a planned node and real-world hardware node are associated in the system, and configurations can be deployed to the physical node as part of a fabric intent.

In addition to the overall inventory, each fabric intent has a dedicated inventory that is specific to that one fabric intent. The nodes in the fabric intent inventory list can include planned nodes and real-world hardware. This list is a subset of the full inventory maintained by the system; it shows only the nodes participating in the selected fabric intent.

From these inventories, you can view specific details about both the planned nodes and real hardware in your fabrics. The columns of the inventory page show specific information about each inventory item, such as status, serial numbers, MAC addresses, and so on.

From the **Inventory** menu, you can access several methods to associate any planned nodes to real-world hardware.

### 9.1 Viewing the overall inventory

#### About this task

There are two types of inventories: the overall inventory and fabric intent-specific inventories. You can view either the overall inventory of nodes, which is a complete inventory of nodes known to the system, or a subset of these nodes which are contained in a specific fabric intent.

Follow this procedure to open the overall inventory known to the system.

#### Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** In the main menu, select  **Inventory**.

You can filter and sort the inventory list according to specific parameters. See [Lists](#) for details on searching and sorting.

#### Expected outcome

The inventory view opens, showing the complete set of nodes in the system's database.

## 9.2 Viewing the inventory of a single fabric intent

### About this task

There are two types of inventories: the overall inventory and fabric intent-specific inventories. You can view either the overall inventory of nodes, which is a complete inventory of nodes known to the system, or a subset of these nodes which are contained in a specific fabric intent.

Follow this procedure to open the inventory associated with a single fabric intent.

### Procedure

**Step 1.** Open a fabric intent in the **Fabric Design** view as described in [Viewing a fabric intent](#).

**Step 2.** In the **View** drop-down, select **Fabric Inventory**.

#### Expected outcome

A list of all nodes participating in the current fabric intent displays.

**Step 3.** To return to the **Fabric Design** view, select **Fabric Design** in the **View** drop-down.

You can filter and sort the fabric intent inventory list according to specific parameters. See [Lists](#) for details on searching and sorting.

## 9.3 Viewing more information about a node in the inventory

### About this task

Follow this procedure to view more information about an individual node in the inventory list.

### Procedure

**Step 1.** With the inventory view open, select a node in the displayed list.

**Step 2.** Click ⓘ at the top of the right panel.

Figure 34: Information for a node in the inventory



## 9.4 Viewing node states

Nodes in the inventory appear in various possible states. With the inventory view open, the State column displays the states of each node in the inventory. The possible states are:

<b>Booted</b>	Real hardware booted in the Fabric Services System. The node is not yet associated with a planned node.
<b>In Discovery</b>	Association of an individual planned node to real hardware has started, and is waiting for the connection.
<b>Ready</b>	An individual planned node has associated successfully with real hardware.
<b>Not Ready</b>	An individual planned node has not successfully associated with real hardware. When a node is in the "Not Ready" state, error codes describe the reason the association was unsuccessful. The following error codes are displayed in the Detailed Status column:

- **Unavailable:** the connection failed.
- **ChassisMismatch:** the planned node's chassis type is different from the real hardware.
- **VersionMismatch:** the planned node's image version is different from the real hardware.

**Unassociated**

A planned node is not associated with any real hardware. The planned node may have made an association using a serial number, but the ZTP process has not completed.

## 9.5 Inventory manipulation

From the Inventory, you can manage the nodes known to the system. You can modify the basic information about individual nodes, associate any planned nodes to real-world hardware, and manage items by importing and exporting information using a spreadsheet.

Each time you change the device associations for a specific node, the notifications menu indicates a Device Status Change update for each node that has been updated.

### 9.5.1 Editing node information

**About this task**

You can edit the serial number and description information of individual real nodes in the inventory. For Digital Sandbox nodes, you can only edit the description. Follow this procedure to edit a node.

**Procedure**

- Step 1.** Open the inventory by following one of the procedures in [Viewing the overall inventory](#).
- Step 2.** At the right edge of the row, click  to open the inventory menu.
- Step 3.** Select **Open** from the displayed action list.
- Step 4.** Either enter a serial number or edit the existing serial number.  
To edit the serial number, double-click the **Serial Number** field. Optionally, you can edit the description.  
For real-world hardware and previously associated nodes, you can also edit the **RackID**.
- Step 5.** Click **SAVE** to save your changes.

### 9.5.2 Viewing the configuration file for a single node

**About this task**

You can view the configuration file that the system has generated for each node in the inventory. Viewing the configuration file can be helpful for verifying the precise configuration that is planned for the node, and possibly revising the configuration if needed.

Follow this procedure to view the current configuration planned for a single node.

## Procedure

- Step 1.** Open the inventory to see the complete list of nodes.
- Step 2.** Find the row for the node in question. At the right edge of the row, click  to open the inventory menu.
- Step 3.** Select **Inspect Configuration** from the displayed actions list.
- Step 4.** Optionally, do any of the following:
  - To save the fabric configuration in a local file, click  and choose a location and name for the file.
  - To copy a portion of the fabric configuration, select the portion and click . The selection is added to your clipboard.
  - To find a particular string of text within the fabric configuration, click  and enter the text string. The first instance is highlighted; use the arrows to navigate forward or backward to additional instances, or click ALL to highlight all instances simultaneously.
- Step 5.** Click the  at the upper right of the overlay to return to the list of nodes.

### 9.5.3 Associating individual planned nodes with real-world hardware

When designing a fabric intent, you can incorporate planned nodes into the design. You can use the Inventory view to associate individual, pending nodes with their real-world counterparts. For 100% fabric intent deployment, the planned node should be associated with real-world hardware. However, each planned node does not need to be associated with real-world hardware to deploy the 0% fabric intent.

When you associate planned nodes and real-world hardware, and generate a fabric intent containing those nodes, the system performs the association between the planned and real devices via the node discovery and ZTP process. A configuration file is generated containing the association details (see [Node discovery](#)).

#### 9.5.3.1 Associate a planned node

##### About this task

Follow this procedure to associate an individual planned node with its real-world counterpart. Alternatively, you can bulk-associate a series of nodes using a spreadsheet. See [Uploading data about a series of inventory items](#).

##### Procedure

- Step 1.** Open the inventory by following one of the procedures in [Viewing the overall inventory](#).
- Step 2.** Find the row for the planned node in question. At the right edge of the row, click  to open the inventory menu.  
If necessary, use the controls at the top of each column to filter the list as described in [Lists](#).
- Step 3.** Do one of the following:
  - Select **Open** from the displayed actions list. In the Update Inventory dialog, update the **Serial Number** field with the serial number of the real hardware item, then click **SAVE**.

- Select **Associate** from the displayed actions list. In the Associate Real Device dialog, select one of the displayed devices, then click **ASSOCIATE**.

**Expected outcome**

The system associates the planned nodes with the real-world hardware. In the UI, the status of the node is set to In Discovery, then when the process completes, the status of the node is set to Ready.

### 9.5.3.2 Associate using the Serial Number column

**About this task**

Follow this procedure to associate an individual planned node with its real-world counterpart using the **Serial Number** column. Use this simple method to quickly add a new serial number without opening a series of menus in the UI.

Alternatively, you can bulk-associate a series of nodes using a spreadsheet. See [Uploading data about a series of inventory items](#).

**Procedure**

**Step 1.** Open the inventory by following one of the procedures in [Viewing the overall inventory](#).

**Step 2.** Select a row corresponding to one of the planned nodes and do the following:

- a. Click in the **Serial Number** column for that row.
- b. Enter the serial number for the real-world node that corresponds to this planned node.
- c. Wait for the real-world node to adopt a Ready state.

**Expected outcome**

Before the association, the planned node and real-world hardware each appeared on individual rows in the inventory. After the association completes, the planned and real-world hardware combine on the same row, as they are now functional nodes in the fabric.

### 9.5.4 Disassociating planned nodes from real hardware

**About this task**

If a planned node is associated with a real-world hardware device, and you want to change this association to a different real-world device, you must remove the existing association. If the node is not currently deployed in a fabric intent, the Inventory allows you to perform the disassociation.

You cannot disassociate a planned node from real-world hardware while the node is in the In Discovery state; the **Disassociate** action is disabled.

When creating a candidate fabric intent, you can disassociate a planned node from real-world hardware while the node is in the Ready state; however, the node loses all established end service bindings with the real hardware.

If the associated node is deployed as part of a fabric intent, but has not been successfully connected to real-world hardware (that is, the node has never been in the Ready state) you can disassociate it. If you want to replace a node that is part of the deployed fabric intent, you must perform a maintenance operation. See [Maintenance intents](#) for details on creating a node replacement maintenance intent.

Follow this procedure to disassociate a planned node from real hardware.

## Procedure

**Step 1.** Open the inventory by following one of the procedures in [Viewing the overall inventory](#).

**Step 2.** Find the row for the planned node in question. At the right edge of the row, click  to open the inventory menu.

If necessary, use the controls at the top of each column to filter the list as described in [Lists](#).

**Step 3.** Select **Disassociate** from the displayed actions list.

**Step 4.** Click **OK**.

### Expected outcome

The system disassociates the planned and real nodes.

## What to do next

After you disassociate a real-world hardware device from a planned node, the real-world hardware item remains in the inventory. It can be re-associated with the planned node, as described in [Associating individual planned nodes with real-world hardware](#). However, if you would like to associate the real-world hardware device to a different planned node, you must update the real-world hardware device's configuration; see [Changing node associations](#).

When a fabric intent is deployed, if any previous versions of that deployed intent included unassociated real-world hardware items, those hardware items are removed. When a fabric intent is deleted, any associated real-world hardware in that intent is removed from the inventory.

## 9.5.5 Changing node associations

When you disassociate a successfully connected planned node from a real-world hardware device (that is, the node has been in the Ready state), you can only continue to use the planned node and hardware device if the hardware device is again reassociated to the same planned node. During the initial association, a configuration file is generated by the system containing configuration information, and the node is discovered by the initial ZTP process. This configuration is unique to the real-world hardware; you cannot associate the planned node with a different real-world hardware device because the certificates will not match. In this scenario, to change the node association, you must manually trigger the ZTP process for the real-world hardware in the management stack so that it receives a new configuration and updated software image. You can then perform a new association.

For example, if you would like to change the association of a planned node from one real-world hardware device (leaf-1), to a secondary real-world hardware device (leaf-2), you will be unsuccessful trying to complete this in the system's UI. This is true even if the hardware devices (leaf-1 and leaf-2) are the same type (that is, both could be 7220 IXR-D3 chassis). Each node has a unique serial number and system name, which is specified in the configuration file. If you generated a certificate on a node for an initial association, then disassociate the node, and try to re-associate with a different node, the action will fail as the certificate provided was intended for the original node.

This scenario is true if you use either the Fabric Services System DHCP server, or an external DHCP server.

### 9.5.5.1 Node associations in candidate fabric intents

If you create a new candidate version of a deployed fabric intent, you can configure new associations between planned nodes and real-world hardware that differ from the associations in the initial version.

The system does not enforce any associations established in the initial version of the intent. Rather, in the candidate version, you can disassociate a previously successful connection between a deployed node and real hardware, resulting in loss of connection and services for the deployed nodes.

When creating a candidate fabric intent, the system assigns an incremented version number; see [Creating a new version of a fabric intent](#). Within the candidate version, you can revise the associations as described in [Associating individual planned nodes with real-world hardware](#) and [Disassociating planned nodes from real hardware](#).

Any revised node associations in the candidate version of the fabric intent appear in the dedicated fabric intent inventory and the overall inventory. When you associate a planned node, the association is versionless. For this reason, if you discard the candidate version of the intent, the node associations do not change.

## 9.5.6 Updating the system name of a node in a fabric inventory

### About this task

You can update the system name of a single node in a fabric intent inventory. To make a bulk-update to a series of nodes, see [Modifying the system name of nodes in an existing fabric intent](#). System names cannot be updated from the overall inventory. You must open a fabric intent and update the node's system name from the fabric intent inventory.

You cannot update the system name of a node if it is already associated with real-world hardware or is in the Deployed state.

Follow this procedure to update the system name of a node in a fabric inventory.

### Procedure

- Step 1.** Open the fabric intent inventory as described in [Viewing the overall inventory](#).
- Step 2.** For the node you want to update, at the right edge of the row, click  to open the inventory menu.
- Step 3.** Select **Update SysName** from the displayed actions list.
- Step 4.** In the editable field, specify the new system name for the node.
- Step 5.** Click **SAVE**.

## 9.5.7 Viewing platform details for nodes in the inventory

### About this task

In the **Inventory** view, you can view specific platform details about each node and the real-world hardware that it is associated with.

For deployed nodes, you can open the **Platform Details** dialog to display information about the chassis components, including fans, power supply units (PSUs), control processing modules (CPMs), and linecards. Specifically, you can view the operational state for each chassis component, and view if there are any current temperature alarms for that component. When the system raises an alarm for any of these components, the platform details shows the precise inventory objects that are affected.

Follow this procedure to view specific platform details on a node in the inventory.

## Procedure

- Step 1.** Open the fabric intent inventory as described in [Viewing the overall inventory](#).
- Step 2.** For the node you want to view, at the right edge of the row, click  to open the inventory menu.
- Step 3.** Select **Show Platform** from the displayed actions list.



**Note:** The **Show Platform** action is only selectable if the planned node is associated with real hardware and the node is deployed in a fabric intent.

- Step 4.** Observe the details in the dialog. You can view the operational state and if there are any current temperature alarms for the component.
- Step 5.** Click either the **CLOSE** button or the  at the top of the dialog to return to the list of nodes.

## What to do next

See [Alarms](#) for more information about how to create and define specific system alarms.

## 9.5.8 Exporting data about inventory items

You can export the information from the inventory into a spreadsheet (.csv) format. The spreadsheet lists either a complete or partial inventory of devices in the network (depending on the selections you make before you export). The exported fabric information shows a list of device names and the serial numbers associated with those devices. In a fabric intent inventory, you can also export a spreadsheet that lists the mapping of system names in the fabric.

The inventory can contain a combination of both actual equipment and pre-planned equipment that is not yet purchased or incorporated into the network.

A fabric inventory spreadsheet can be used when planning and implementing a network. It can be distributed to personnel that do not have access to Fabric Services System. For example, a spreadsheet can be used for device procurement, internal asset management, or can be distributed to DC contractors that perform physical cabling to build the network.

You can export fabric data from either the overall inventory of fabric items, or the fabric intent inventory of a specific fabric intent.

### 9.5.8.1 Exporting an overall inventory spreadsheet

#### About this task

In the overall inventory of fabric items, you can export a spreadsheet containing all of the information listed in the Inventory. The spreadsheet can be distributed to personnel that do not have access to the Fabric Services System UI.

Follow this procedure to export the overall inventory.

#### Procedure

- Step 1.** Open the overall inventory by following the procedure in [Viewing the overall inventory](#).
- Step 2.** From the inventory, click the  menu at the upper right of the page.
- Step 3.** Select **Export**, then **CSV export** from the resulting list.

**Step 4.** When the download completes, either open the .csv file or save it to your local system.

**Expected outcome**

The resulting spreadsheet displays all of the data shown on the **Inventory** screen.

### 9.5.8.2 Exporting association data from the overall inventory

#### About this task

In the overall inventory of fabric items, you can download an associate mapping file to export the details about the inventory items.

Follow this procedure to export multiple items from the overall inventory to a spreadsheet.

#### Procedure

**Step 1.** Open the inventory by following the procedure in [Viewing the overall inventory](#).

**Step 2.** Select two or more rows by checking the box at left side of each row.

**Step 3.** Click the **Download Associate Mapping File** icon  .

**Step 4.** When the download completes, either open the .csv file or save it to your local system.

**Expected outcome**

The resulting spreadsheet displays information about each inventory item you selected. The **name** column displays the name of the item and the **serialNumber** column displays the associated serial number (if associated).

### 9.5.8.3 Exporting association data from a fabric intent inventory

#### About this task

In the inventory of a specific fabric intent, you can download an associate mapping file or a system name mapping file to export the details on the inventory items contained within the fabric intent.

Follow this procedure to export multiple items from the fabric intent inventory of a specific fabric intent to a spreadsheet.

#### Procedure

**Step 1.** Open the inventory by following one of the procedures in [Viewing the overall inventory](#).

**Step 2.** Select two or more rows by checking the box at left side of each row.

**Step 3.** Click either the **Download Associate Mapping File** icon  or **Download SysName Mapping File** icon .

**Step 4.** When the download completes, either open the .csv file or save it to your local system.

**Expected outcome**

The resulting spreadsheet displays information about the multiple selected items.

If you downloaded the Associate Mapping File, the spreadsheet **name** column displays the name of the items and the **serialNumber** column displays the associated serial numbers (if associated). If you downloaded the SysName Mapping File, the spreadsheet **name** column displays the name of the nodes and the **newSystemName** column displays a blank field where you can add a new system name.

## 9.5.9 Uploading data about a series of inventory items

You can upload (import) information from a spreadsheet (.csv) format to the Fabric Services System. The spreadsheet contains information to complete the missing details of planned items in the Inventory, or update details on existing items. You can upload a spreadsheet of details related to one or more planned fabric items.

After you have exported an inventory spreadsheet (see either [Exporting association data from the overall inventory](#) or [Exporting association data from a fabric intent inventory](#) procedures) and updated it with new information (such as serial numbers to associate specific real-world hardware to each pending item in the inventory), you can import the spreadsheet back into the system to bulk-associate pending nodes with real-world counterparts.

Import fabric data from either the overall inventory of fabric items, or the fabric intent inventory of a specific intent.

### 9.5.9.1 Uploading association data to the overall inventory

#### About this task

Follow this procedure to upload data about new inventory items installed on the network.

#### Procedure

- Step 1.** Open the inventory by following the procedure in [Viewing the overall inventory](#).
- Step 2.** Click **UPLOAD ASSOCIATE MAPPING FILE**.
- Step 3.** Select the spreadsheet from your local system.
- Step 4.** In the selection window, click **Open**.

### 9.5.9.2 Uploading association data to the fabric intent inventory

#### About this task

Follow this procedure to upload data about new inventory items for a specific fabric intent.

#### Procedure

- Step 1.** Open the fabric intent inventory by following the procedure in [Viewing the overall inventory](#).
- Step 2.** Click **UPLOAD ASSOCIATE MAPPING FILE**.
- Step 3.** Select the spreadsheet from your local system.
- Step 4.** In the selection window, click **Open**.

### 9.5.9.3 Modifying the system name of nodes in an existing fabric intent

#### About this task

If you need to change the system name of items in an existing fabric intent, you can download the System Name Mapping file, update the information, then upload the file to add the node names to the fabric intent.

The System Name Mapping file allows you to bulk-update the nodes the inventory instead of updating each individually.

Follow this procedure to modify system name mapping data for a specific fabric intent.

### Procedure

- Step 1.** Open the fabric intent and download the System Name Mapping file (.csv) by following the procedure in [Exporting association data from a fabric intent inventory](#).
- Step 2.** Open the downloaded file on your local system.  
The spreadsheet displays two columns. The **name** column displays the current names of the nodes and the **newSystemName** column displays a blank field where you can add the new system names.
- Step 3.** In the **newSystemName** column, enter a new name for each node you would like to change.
- Step 4.** Save the updated spreadsheet on your local system.
- Step 5.** Click **UPLOAD SYSNAME MAPPING FILE** and select the updated spreadsheet from your local system to upload the changes.

### Expected outcome

The names of the nodes in the inventory are updated as specified in the spreadsheet. You can verify the changes immediately in the **Fabric Inventory** view.

## 10 Alarms

In the Fabric Services System, alarms arise when a managed object enters an undesirable state; for example, if a managed node goes out of service, an associated alarm is raised.

The Fabric Services System supports the following alarm types, each with its own subtypes:

- Communication
- Configuration
- Environment
- Equipment
- Operational

The Fabric Services System includes the following tools that can help you manage alarms:

- the Alarms panel on the Dashboard, which summarizes current alarms; for more information about this panel see [The dashboard](#)
- the Alarm List, which you can use to view and manage individual alarms
- the Policy manager, which you can use to customize the severity level for specific types of alarm, or to suppress alarms of a specific type entirely.

For a list of individual alarms currently supported by the Fabric Services System, see [Appendix: Supported alarms](#).

### Alarm states

In the Fabric Services System, an alarm can adopt the following states:

- **Acknowledged:** An acknowledged alarm still displays in the Alarm List; but when viewing details for the individual alarm its state displays as Acknowledged and any note you added to the alarm while acknowledging it is displayed as well. You can use the Acknowledge state as the basis for filtering or sorting the alarm list.
- **Closed:** A Closed alarm still displays in the Alarm List, but this state can be the basis for filtering the alarms included in the list. Closing an alarm does not resolve the condition that caused the alarm to be raised in the first place.
- **Cleared:** An alarm is Cleared when the condition that raised the alarm has been resolved. Unlike Acknowledged and Closed, the Cleared state cannot be assigned manually by a Fabric Services System operator. Only the device or devices that raised the original alarm can determine and communicate its closure.

### 10.1 Alarm list

#### About this task

The Alarm List displays a list of current alarms known to the Fabric Services System. From this page you can view details about each alarm's state, and also acknowledge any alarms.

To view and manage alarms with the Alarm List:

## Procedure

**Step 1.** From the main menu, select  **Alarm List**.

### Expected outcome

The alarm list displays, showing all active alarms (where "active" refers to alarms that have not been cleared).



**Note:** Cleared alarms are not included in this list because the "Cleared" filter is set to "False" by default. To view cleared alarms in this list, clear that filter. For more information about managing lists and filters, see [Lists](#).



**Note:** A set of default columns display in the alarms list:

- Severity
- Alarm type
- Node name
- Resource name
- Cleared
- Occurrence
- Last Raised

There are other columns available to show more information about each alarm. For information about adding or removing columns from any list, see [Lists](#).

**Step 2.** To view details about an alarm and its state:

- Select an alarm in the list.
- At the right edge of the row, click  and select **State Details** from the displayed action list.
- Click the **ALARM STATE** tab to view details about the alarm's severity, a description of the alarm, and the time it was raised.
- Click the **OPERATOR STATE** tab to view the state assigned by the operator to address the alarm (either Acknowledged or Closed).
- When you are finished, click **CLOSE** to return to the Alarm List page.

**Step 3.** To acknowledge an alarm:

Acknowledging an alarm marks it as received, but does not clear the alarm from the Alarm List.

- Select an alarm in the list.
- At the right edge of the row, click  and select **Acknowledge** from the displayed action list.
- Optionally, enter any comments about the acknowledgement in the **Additional Info** field.
- Click **SAVE**.

### Expected outcome

The Alarm is marked as Acknowledged (but not Closed).

**Step 4.** To Close an alarm:

Closing an alarm prevents the alarm from appearing in the Alarm List, but does not resolve the condition that raised the alarm in the first place.

- a. Select an alarm in the list.
- b. At the right edge of the row, click  and select **Close** from the displayed action list.
- c. Optionally, enter any comments about the acknowledgement in the **Additional Info** field.



**Note:** The text you enter here is displayed in the **Additional Info** column of the **OPERATOR STATE** tab in the Alarm Details overlay.

- d. Click **SAVE**.

## 10.2 Policies

### About this task

Policies allow you to customize the severity level associated with individual supported alarms. A policy affects the alarm types of all future alarms raised; it does not retroactively modify any existing alarms of the same type.

Each policy can include a start time and an end time; these are boundaries on the time of day during which the policy applies. An alarm raised outside these boundaries has its default severity instead of the severity level defined by the policy. If no start and end time are defined, the policy is always active.

You can also use a policy to suppress an alarm entirely while the policy is in effect.

The Fabric Services System supports the definition of a policy's scope in two mutually exclusive ways:

- by key value, which allows you to trigger a policy based on the name of the object (node, fabric, intent, or region) affected by the alarm
- by alarm category and type, to apply the policy regardless of the object affected.

To customize an alarm's severity:

### Procedure

**Step 1.** From the main menu, select  **Policies**.

**Step 2.** Click **+ CREATE A POLICY**.

**Step 3.** Enter a **Name** and **Description** for the policy.

**Step 4.** In the Policy Definition panel, enter a **Start Time** and **End Time** for the policy.

An alarm that would be affected by this policy uses the customized severity level only if it is raised during this period. If it is raised outside this period, it uses the default severity level.

**Step 5.** Do one of the following:

- To configure a policy based on the object it affects, click the **Key Value** toggle to enable it and go to step 6.
- To configure a policy based on alarm category and type, leave the **Key Value** toggle disabled and go to step 7.

**Step 6.** With the **Key Value** toggle enabled, do the following:

- a. Click on the **Key Value Objects** drop-down and select one or more of the displayed Key Value candidates:
  - Node Name

- Fabric Name
- Intent Name
- Region Name

**Expected outcome**

For each selected item, a field displays.

- Use the displayed field or fields to provide the unique name of each object type you selected. This name identifies the unique object of that type for which the system changes the alarm severity or suppresses alarms, depending on how you configure the remainder of the policy.
- Go to step 8.

**Step 7.** In the Policy Definition panel, do the following:

- Select an **Alarm Category** from among the following:
  - Communication
  - Equipment
  - Operational
  - FSS
- Click the **Alarm Type** drop-down and check the box beside one or more alarms types in the displayed list.
- Select an **Alarm Severity**:
  - Major
  - Minor
  - Critical
  - Warning
  - Default

**Step 8.** Configure the way this policy modifies alarms:

- Select a **Priority** from 0 to 9, with 0 being the highest priority.
- Optionally enable **Suppress Alarms** toggle.

Enabling this option means that alarms of this type are disabled and are not triggered while the policy is in effect.
- Optionally enable the **Deployed Intent Alarms Only** toggle.

This option applies to "Communication – Interface Down" alarms. If this option is enabled, alarms are raised only on interfaces that are part of the intent configuration.

**Step 9.** Click **CREATE**.

# 11 Operations views

To help you assess the overall health of your fabrics, the Fabric Services System includes some high-level perspectives on the current status of all managed regions and fabrics. These views are intended to help you assess at a glance what is working and what is not, what is configured or misconfigured, and to understand what other related objects could be affected by an unhealthy entity.

From these high-level views that begin at the regions level, you can drill down to see the health of individual fabrics, nodes, and links within a region.

Health is a reflection of the operational state of nodes, links, and any node-based alarms within the set of managed fabrics. Although health alarms are distinct from intent deviations and deployment status, they may be related.

You can access the main operational view by selecting  **Operation and Health Insights** from the system's main menu. The resulting page includes two main overlays:

-  **Health**, which displays information about objects subject to alarms, and their operational state (Up, Down, or Missing Hardware)
-  **Operational Deviation**, which indicates where there are errors within the physical topology of a fabric.

In each case, this information is conveyed through a combination of graphical display for overall status, and the information panel on the right side of the page that provides more details.

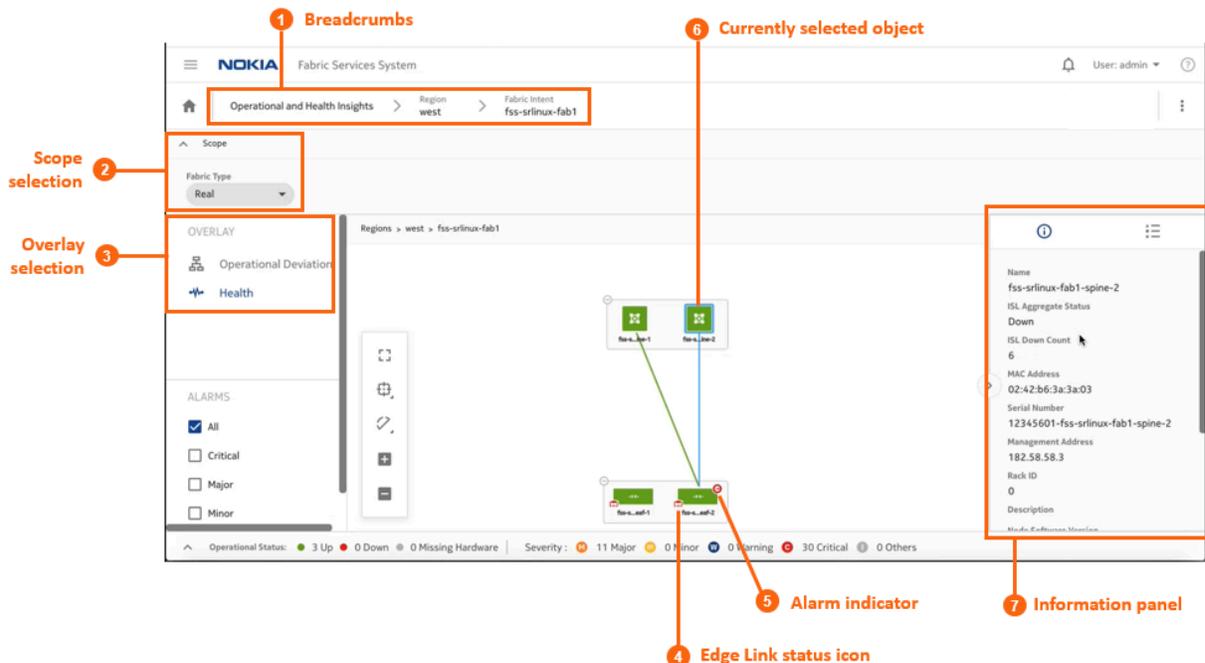


Table 21: Operational health display

#	Description
1	Breadcrumbs: as you drill down into a region, this displays the level of your current view and the sequence of objects between that view and the region level.
2	Scope selection: use this drop-down to control whether the view displays Real objects, or simulated Digital Sandbox objects.
3	Overlay selection: use this panel to select one of the available views, either Health or Operational Deviations.
4	Edge Link status icon: this icon is only present on Leaf nodes, and indicates the presence (and status) of edge links leading to customer assets. Possible statuses are Up (green) or Down (red).
5	Alarm indicator: this icon indicates the aggregate alarm level for an object. If absent, there are no alarms on the object or its child objects. If present, it indicates (by its color and letter) the single highest-level severity of any alarms that are active on the object or any of its children.
6	Currently selected object: a light blue outline indicates the currently selected object within the view. The same light blue is superimposed on any ISLs associated with that object.
7	Information panel: open this panel by clicking the a the right of the page. The panel displays information about the currently selected object. Of particular interest are the aggregate status and count indicators: <ul style="list-style-type: none"> <li>• For a region: Node Aggregate Status and Node Down Count</li> <li>• For a fabric: Node Aggregate Status, Node Down Count, Interface Aggregate Status, Interface Down Count</li> <li>• For spine nodes: ISL Aggregate Status and ISL Down Count</li> <li>• For leaf nodes: ISL Aggregate Status, ISL Down Count, Edge Link Aggregate Status, and Edge Link Down Count.</li> </ul>

### The Health overlay

-  **Health**, which displays information about:
  - which objects that are subject to alarms
  - for parent objects, the single highest-severity alarm on any of its child objects
  - which nodes and links are down
  - for parent objects the proportion of immediate child objects that are up or down
-  **Operational Deviation**, which indicates where there are errors within the physical topology of a fabric. A deviation in this context represents:
  - a port that should be connected to another port, but is not
  - a port that should not be connected to another port, but is
  - a port that should be connected to a specific other port, but is connected to a different port instead

Displays on the Health overlay follow the conventions for other maps in the Fabric Services System UI. The possible states that can be displayed vary depending on whether you are looking at the Health page or the Operational Deviations page.

Table 22: Operational states and alarm severities

Information type	Display				
Operational states (Health)					
	Operational Up	Operational Down	Missing Hardware		
Operational states (Deviations)					
	Intent Matching	Operational Deviation	Missing Hardware		
Alarm severities					
	Major	Minor	Warning	Critical	Other (Information)

### The Operational Deviation overlay

The **Operational Deviation** overlay indicates where there are errors within the physical topology of a fabric. A deviation in this context represents:

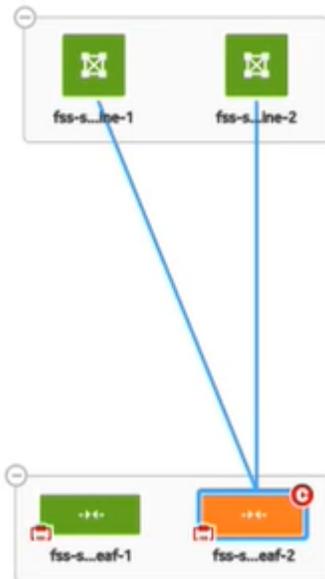
- a port that should be connected to another port, but is not
- a port that should not be connected to another port, but is
- a port that should be connected to a specific other port, but is connected to a different port instead

Any object that is the subject of deviations is shaded orange in the overlay topology.

Figure 35: Operational deviation on a fabric



Figure 36: Operational deviation on a leaf node



## 11.1 Viewing the operational topology

### Procedure

**Step 1.** From the main menu, select **Operational and Health Insights**.

#### Expected outcome

The **Operational and Health Insights** page displays, showing a map of all configured regions. (This is currently limited to a single region.)

**Step 2.** In the **OVERLAY** drop-down, select **Health**.

**Step 3.** Optionally, select the type of fabrics to display by making a selection in the **Fabric Type** drop-down list.

The drop-down supports two fabric types:

- Real (default)
- Digital Sandbox

**Step 4.** To learn more about any region on the map, do the following:

**a.** Check the circle in the center of the region.

If the circle contains any red, that indicates the proportion of fabrics within the region that have operational issues. You can expand the region and view more information about its status by going to sub step 4.b.

**b.** Click the region icon to select it, then click the **i** icon at the right of the page to open the information panel.

The information panel displays more information about the selected region, including the number of fabrics and number of links it contains.

**Step 5.** To learn more about the fabrics within the region, do the following:

a. Expand the region icon.

**Expected outcome**

The region icon expands to show all of the fabrics it contains. Fabrics that contain nodes with operational issues are shaded red.

b. Select a fabric and, if the information panel is not open, click the ⓘ icon at the right of the page to open the information panel.

The information panel displays more information about the selected fabric. Pay particular attention to the following values:

- **Node Aggregate Status:** this summarizes the overall state of nodes within the fabric.
- **Node Down Count:** this is the number of nodes within the fabric that are currently operationally Down.
- **Interface Down Aggregate Status:** this provides some information about the proportion of interfaces within with the fabric that are experiencing operational issues.
- **Interface Down Count:** this is the number of interfaces within the fabric that are experiencing operational issues.

**Step 6.** To navigate to another view to learn more about an individual fabric, right-click the fabric on the map and select one of the following from the resulting contextual menu:

- **Open Fabric Design** to open the fabric intent's Design view
- **Open Fabric Inventory** to open the fabric intent's Inventory view
- **Show Intent Alarms** to open the Alarms List.

**Expected outcome**

The selected view displays, replacing the **Regions and Health** view. The new view adopts the context of the fabric whose row you selected.

**Step 7.** To learn more about the spine nodes within a single fabric, do the following:

a. Double-click the fabric icon on the **Operations and Health Insights** map.

**Expected outcome**

The map shows the topology of the selected fabric, including any backbone, spine, and leaf groups.

b. Select and expand a spine cluster.

**Expected outcome**

The expanded group shows the individual nodes within that cluster. Node icons are shaded green if Up, red if Down, and gray if Unassociated with hardware.

c. Select an individual spine node within the expanded group and, if the information panel is not open, click the ⓘ icon at the right of the page to open the information panel.

The information panel displays more information about the selected node. Pay particular attention to the following values:

- **ISL Aggregate Status:** this provides some information about the proportion of ISLs associated with the selected node that are experiencing operational issues.

- **ISL Down Count:** this is the precise number of ISLs associated with the selected node that are experiencing operational issues.

**Step 8.** To learn more about a leaf nodes within a single fabric, do the following:

a. Select and expand a leaf group.

**Expected outcome**

The expanded group shows the individual nodes within that cluster. Node icons are shaded green if Up, red if Down, and gray if Unassociated with hardware. Similarly, the Edge Link icon at the bottom right of the leaf node icon is shaded green if all edge links are Up, or red if at least some edge links are Down.

b. Select an individual leaf node within the expanded group and, if the information panel is not open, click the ⓘ icon at the right of the page to open the information panel.

The information panel displays more information about the selected node. Pay particular attention to the following values:

- **Edge Link Aggregate Status:** this provides some information about the proportion of edge links associated with the selected node that are experiencing operational issues.
- **Edge Link Down Count:** this is the precise number of edge links associated with the selected node that are experiencing operational issues.
- **ISL Aggregate Status:** this provides some information about the proportion of ISLs associated with the selected node that are experiencing operational issues.
- **ISL Down Count:** this is the precise number of ISLs associated with the selected node that are experiencing operational issues.

## 11.2 Viewing operational insights

### Procedure

**Step 1.** From the main menu, select  **Operational and Health Insights**.

**Step 2.** Optionally, select the type of fabrics to display by making a selection in the **Fabric Type** drop-down.

The drop-down supports two fabric types:

- Real (default)
- Digital Sandbox

**Step 3.** In the **OVERLAY** drop-down, select  **Operational Deviation**.

**Step 4.** Expand the region icon.

**Expected outcome**

The region icon expands to show all of the fabrics it contains. Fabrics that contain nodes that are the subject of deviations are shaded orange.

**Step 5.** Double-click a fabric icon on the **Operations and Health Insights** map.

**Expected outcome**

The expanded fabric cluster shows the individual nodes within that cluster. Nodes that are the subject of deviations are shaded orange.

## 12 Digital Sandbox

The Fabric Services System Digital Sandbox is a network simulator that can emulate data center fabric designs ("underlays") and the workload constraints configured upon those fabrics ("overlays").

The Digital Sandbox normally runs on a three-node Kubernetes cluster. Each SR Linux node is emulated as its own virtual machine within the cluster, running its own copy of the SR Linux operating system like the real node it represents.

Before you can use the Fabric Services System Digital Sandbox, you must install its software components and perform any configuration steps described in the *Fabric Services System Software Installation Guide*. This ensures that the Digital Sandbox software is ready to simulate your fabrics and workloads, and is ready to communicate with the Fabric Services System to receive model data and send status updates.

In its current form, the Digital Sandbox can emulate a region, the structures of fabrics within that region, and the workload constraints that are configured upon those fabrics (including the edge links that are referred to by the workload). It does not yet simulate dynamic features like traffic flow between the simulated nodes and their endpoints.

The Digital Sandbox requires its own license, purchased separately from the license for the Fabric Services System itself.

### 12.1 Integration with the Fabric Services System

Most interaction with the Digital Sandbox is performed using CLI commands and API calls.

In integrated mode, the Digital Sandbox can communicate with the Fabric Services System, receiving configuration data for fabric and workload designs, and returning status updates for those intents. The Fabric Services System UI does not support all of the available Digital Sandbox operations; it is only used to design fabric intents and workload VPN intents, and to send those configurations to the Digital Sandbox for further action.

A technically proficient user who is familiar and experienced with the Digital Sandbox CLI could configure fabrics, workloads, and participating endpoints using only the Digital Sandbox CLI or REST API calls. But for most operators, taking advantage of integration with the Fabric Services System makes these operations much faster and easier.

#### 12.1.1 Digital Sandbox status display

When using the integrated mode, the Fabric Services System UI displays the status of the Digital Sandbox in the lower-left corner of the fabric intents geographical map. Possible statuses are:

- Unavailable: the Digital Sandbox has either:
  - not been installed
  - not been configured for integration with the Fabric Services System
  - been misconfigured, such that the Fabric Services System is pointing to the wrong location for the Digital Sandbox

- has been installed and configured for integration but has not been started
- Running: the Digital Sandbox is installed, started, and ready to receive data.
- Busy: the Digital Sandbox is installed and started, but is processing data recently sent from the Fabric Services System for incorporation into its simulation.

## 12.2 Creating a region

When using the integrated mode, you must create a deployment region from the Fabric Service System's **Regions** page before you can create any fabric intents or workload VPN intents that are destined for the Digital Sandbox.

Create a region by following the procedure [Creating a region](#).

After you have created the region, the Digital Sandbox creates a set of internal structures ("pods") for use in its simulation.

As the Digital Sandbox creates these structures, its status advances from Unavailable to Busy to Running. When it reaches the Running state, you can proceed to create fabric intents and workload intents that are destined for the Digital Sandbox.

### 12.2.1 Modifying a region

#### About this task

If you modify any of the properties of a region, including changes to the fabric and workload intents in the Fabric Services System UI destined for digital sandbox, you must explicitly trigger a corresponding update in the Digital Sandbox.

To update the Digital Sandbox after modifying a region:

#### Procedure

**Step 1.** If the region page is not already open:

- Click  to open the main menu.
- In the main menu, select  **Deployment Regions**. The Deployment Regions page opens, showing a graphical representation of regions already created.

**Step 2.** Right-click the region icon on the Deployment Regions page, and select Update Digital Sandbox from the displayed menu.

#### Expected outcome

In the lower left corner of the page:

- the Digital Sandbox status advances to the Busy state.
- the Digital Sandbox status returns to the Running state.

When the Digital Sandbox has returned to the Running state, the update to its model of the region is complete.



**Note:** If an error occurs during the update, an error indication appears on the lower left corner. Hover over the error indication to display the description of the error in digital sandbox.

## 12.3 Fabric intents and the Digital Sandbox

Working with a fabric intent that is destined for the Digital Sandbox is just like working with a fabric intent destined for real hardware, with the following exceptions:

- When you create the fabric intent, set the Fabric Type to Digital Sandbox instead of Real.
- The deployment threshold for a Digital Sandbox fabric intent is always 100% (whereas the threshold is always 0% for Real fabric intents).
- After you design a Digital Sandbox fabric intent, and before you deploy that intent, you must manually send the updated configuration to the Digital Sandbox using the Update Digital Sandbox command.

After you update the Digital Sandbox and deploy your fabric intent, the Digital Sandbox creates a corresponding set of simulated nodes, their configurations, and their relationships in its own data model.

The representation of the fabric at this stage is coarse; it is limited to the nodes themselves. Endpoints and other lower-level details are not added to the model until subsequently required by workload intents.

### 12.3.1 Creating a fabric intent

#### Prerequisites

Before you create a fabric intent that is destined for the Digital Sandbox, make sure you have followed the procedure for [Creating a region](#).

#### About this task

The procedure to design a fabric intent destined for the Digital Sandbox is nearly identical to designing an intent for deployment to real hardware.

To create a fabric intent that is destined for the Digital Sandbox:

#### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** In the main menu, select  **Fabric Intents**. The Fabric Intents page opens.
- Step 3.** Click the **+ CREATE A FABRIC INTENT** button.
- Step 4.** Select any of the following templates:
  - Flexible Leaf Spine
  - Backbone
- Step 5.** Click the **CREATE** button.

#### Expected outcome

The New Fabric Intents page displays.

At the top of the left-side panel, the name of the template you selected in step 4 is displayed as the "Reference Template".

**Step 6.** In the **Fabric Type** drop-down list, select Digital Sandbox.

**Step 7.** From here, continue with one of the procedures in [Fabric intents](#), depending on the type of fabric you are designing:

- [Creating a new fabric intent: flexible leaf/spine](#)
- [Creating a new fabric intent: backbone](#)

The steps are identical, except that the Deployment Threshold value is fixed at 100% (whereas for a Real fabric intent, this value is fixed at 0%).

## 12.3.2 Updating the Digital Sandbox

### About this task

After you design your fabric intent, but before deploying it, you must update the Digital Sandbox with information about the fabric intent. This causes the Digital Sandbox to create virtual nodes onto which the fabric can be deployed.

Any time you make subsequent changes to the fabric intent, you should follow these steps to again update the Digital Sandbox with the new configuration data.

To update the Digital Sandbox with information about your fabric intent, do the following:

### Procedure

**Step 1.** Open the fabric intent in the **Fabric Design** view, if it is not already open.

**Step 2.** At the upper right of the page click the **More actions** icon (  ) and select **Update Digital Sandbox** from the displayed menu. As a result:

#### Expected outcome

In the lower left corner of the page:

- the Digital Sandbox status advances to the Busy state.
- one by one, each virtual node in the fabric intent advances to the Ready state.
- the Digital Sandbox status returns to the Running state.

When all nodes are in a Ready state and the Digital Sandbox has returned to the Running state, you can deploy the fabric intent.



**Note:** If an error occurs during the update, an error indication appears on the lower left corner. Hover over the error indication to display the description of the error in digital sandbox.

Additional messages can also appear in the lower middle section of the page.

## 12.3.3 Deploying a fabric intent

### About this task

The procedure to deploy a fabric intent to the Digital Sandbox is the same as that for a Real fabric intent:

- add the fabric intent to the deployment pipeline
- from the deployment pipeline, select the fabric intent and click **Deploy**.

### Procedure

For detailed steps, see [Deploying a fabric intent](#).

### Expected outcome

When the fabric has attained the Deployed state, the result in the Digital Sandbox is a new Underlay 1 (UL1) construct, represented by a collection of Kubernetes pods.

At this stage only the nodes themselves are modeled in the Digital Sandbox data. Endpoints, and details about those endpoints such as IP addresses, are not yet present in the model.

## 12.4 Workload intents

There are no differences in workload VPN intent design or deployment when the target is the Digital Sandbox; the procedures are the same as those for workload VPN intents destined for real hardware.

When you design a Workload VPN intent that includes fabrics that were created for the Digital Sandbox, the Fabric Services System sends the workload intent information to the Digital Sandbox for incorporation into its simulation.

### 12.4.1 Creating a workload VPN intent for the Digital Sandbox

Create your workload VPN intent just as described in [Workload VPN intents](#).

When you are finished, you are ready to deploy the workload VPN intent.

### 12.4.2 Deploying a workload VPN intent to the Digital Sandbox

Deploy your workload VPN intent just as described in [Adding a workload intent to the deployment pipeline](#) and [Deploying a workload intent from the deployment pipeline](#).

When you deploy the workload intent, the Digital Sandbox updates the configuration files of the participating, simulated nodes. In the Digital Sandbox, the simulated workload is classified as a Candidate Workload, but is not active; the participating nodes are identified, but endpoint data is not yet present.

If you were to update the workload VPN intent design in the Fabric Services System and re-deploy it, this would entirely overwrite the Candidate Workload's information in the Digital Sandbox. This would be true even if you had updated the workload information in the Digital Sandbox directly using the CLI after the last deployment; the re-deployment would overwrite the workload intent data and erase your changes.

### 12.4.3 Updating the Digital Sandbox

#### About this task

After you deploy your workload intent, you must explicitly update the Digital Sandbox with information about the workload intent. This causes the Digital Sandbox to add endpoint data to the workload model, and the result is an Active Workload.

Any time you make subsequent changes to the workload intent, you should follow these steps to again update the Digital Sandbox with the new configuration data.

To update the Digital Sandbox with information about your fabric intent, do the following:

### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** In the main menu, select  **Deployment Regions**.
- Step 3.** Right-click the region object on the **Deployment Region** map.
- Step 4.** From the contextual menu, select **Update Digital Sandbox**.
- Step 5.** Click **OK**.

### Expected outcome

The Digital Sandbox updates its model based on the latest data in the Fabric Services System. While it does this, the Digital Sandbox status advances from the Running state to the Busy state, and then back to the Running state.

When the Digital Sandbox has returned to the Running state, its simulated model of the workload VPN intent includes all of the participating endpoints.

If you were to update the workload VPN intent design in the Fabric Services System, re-deploy it, and re-update the Digital Sandbox, this would entirely overwrite the Active Workload's information in the Digital Sandbox. This would be true even if you had updated the workload information directly in the Digital Sandbox using CLI after the last update; re-updating would overwrite the workload intent data and erase your changes.

## 13 System administration

System administration encompasses a range of activities that configure or otherwise manage the Fabric Services System application itself.

Some of these operations are performed within the Fabric Services System user interface. Others are performed as command-line operations on the server that hosts the application.

### 13.1 Application settings

The Fabric Service System's **Settings** page provides access to several configuration settings that affect the way the system behaves.

- the Image Management and Software Catalog panels pertain to the SR Linux software images available on the Fabric Services System's file server.
- the Global Configuration panel manages FTP and NTP settings on the SR Linux nodes configured by the system.
- the Common Application Settings page determines the background map image used with specific views, and whether you are using Kibana and Elasticsearch to view log data generated by the system's internal microservices.

#### 13.1.1 Software and image catalogs

##### About this task

The Fabric Services System maintains a catalog of software images which are used to configure nodes within managed fabrics.

The software images in the catalog are used when:

- creating a fabric intent; you must select one version of the SR Linux software which is deployed to all nodes participating in the fabric as part of their configuration. See [Fabric intents](#).
- creating a maintenance intent for changes to a node's software version; you must select one version of the SR Linux software from the catalog which is deployed to all nodes that are subjects of the maintenance intent. See [Maintenance intents](#).

To view the system's current software catalog:

##### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** Select  **Settings**. The Settings page displays.
- Step 3.** Click **Software Catalog**. The current software catalog displays, identifying for each entry:
  - the **Vendor**: currently only Nokia is supported.
  - the **Operating System**: currently only SR Linux is supported.

- the **Software Version**: the system is pre-loaded with selected versions of the SR Linux software image, and this list grows as you load additional versions over time.

### 13.1.1.1 Adding a new software image

#### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** Select  **Settings**. The Settings page displays.
- Step 3.** Click **Image Management**. The current image catalog displays.
- Step 4.** Click **+ADD IMAGE**. The Image Details dialog displays.
- Step 5.** Enter information for the new software image:
  - the **Image URL**: a location from which the system can download the image file to the file server.
  - the **MD5 URL**: a location (on the same file server as the Image URL) from which the system can download the MD5 file to the file server.
  - a **Username** and **Password**: credentials required to download files from the Image URL.
  - Optionally, a **Description** for this software image.
- Step 6.** Click **SAVE**.

#### Expected outcome

The system adds the image to the image catalog and downloads the specified image from the source you provided to its integrated file server.

In addition to the information you provided in step 5, the image catalog displays for each software image:

- **Image Name**: the name of the image file on the file server.
- **MD5**: either True or False, indicating whether the MD5 file has been successfully added.
- **Upload Status**: the image status, which can be any of the following:
  - Done
  - Failed
  - Adding
  - Overwriting
- **Path**: for a successfully obtained image, the path to the image and MD5 files on the integrated file sever.

### 13.1.2 Global configuration settings

Global configuration settings represent data that the Fabric Services System incorporates into the configuration files of any node included within a fabric intent.

### 13.1.2.1 FTP

#### Prerequisites

#### About this task

The FTP toggle enables or disables the use of File Transfer Protocol (FTP) by the nodes configured by the Fabric Services System.

The UI supports two settings for FTP:

- Enabled (toggle ON): nodes configured by the Fabric Services System include a setting to enable FTP on the node.
- Disabled (toggle OFF): nodes configured by the Fabric Services System include a setting to disable FTP on the node.

To configure the FTP setting:

#### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** Select **Settings** .
- Step 3.** Click **Global Configuration**.
- Step 4.** Set the **Enable FTP** toggle to either disabled (left, gray) or enabled (right, blue).
- Step 5.** Click the **SAVE** button at the lower right of the page.

### 13.1.2.2 NTP

#### About this task

The NTP setting specifies one or more Network Time Protocol (NTP) servers; this information is incorporated into the configuration file for all nodes configured by the Fabric Services System. Adding one or more NTP server IP addresses enables the system NTP client on the node, and is used for system clock synchronization.

To add an NTP server that is included in node configurations going forward:

#### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** Select **Settings** .
- Step 3.** Click **Global Configuration**.
- Step 4.** Click **+ ADD**.
- Step 5.** Enter the IP address of an NTP server.  
This should be an ipv4 address with an optional CIDR.
- Step 6.** Click the **ADD** button.
- Step 7.** Repeat steps 4 through 6 until you have added all the required NTP server addresses.
- Step 8.** Click the **SAVE** button at the lower right of the page.

### 13.1.3 Common application settings

The Common Application Settings panel determines the background map image used with certain views.

#### 13.1.3.1 Geomap

##### About this task

The background image displayed for the region map is determined by the system's Geomap Tile Server setting.

To assign or change the background map:

##### Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** Select **Settings**.

**Step 3.** Click **Common Application Settings**.

**Step 4.** In the Geomap panel, enter:

**a. Tile Server:** enter a URL that points to a specific tile server and an image on that server to function as the background for the region map

**b. Attribution:** enter an attribution for the background image.

The attribution field is mandatory.

**Step 5.** Click the **SAVE** button at the lower right of the page.

##### Expected outcome

The selected image displays the next time you open the region map.

## 13.2 User and resource management

The user and resource management menu allows administrators to manage users and their access to resources in Fabric Services System.

Administrators can configure four types of elements related to users and resources:

<b>Roles</b>	Specifies which network resources users or associated user group members can access. You assign network resource access to roles through resource groups.
<b>Resource groups</b>	Specifies resources that users can access. When you create a role, you associate specific resource groups to it. Then, when you assign a role to a user group, the members of the user group gain access to the resource groups associated with the role. For example, resource groups can represent the ability to configure fabrics or QoS policies.
<b>User groups</b>	A collection of users organized according to the type of network activities they are meant to perform. You assign resource access rights to user groups

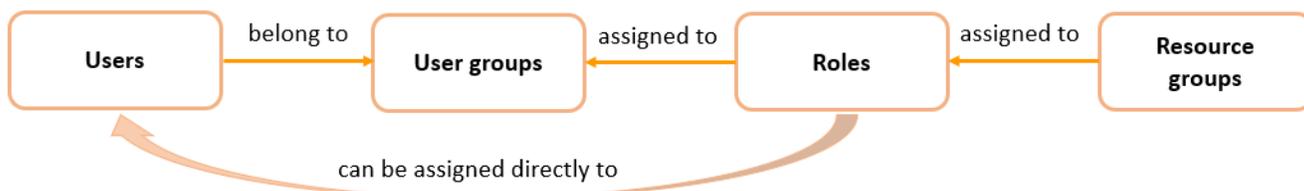
through user roles. When you assign a role to a user group, all access rights defined in the role are inherited by the users of the group.

### Users

Individuals with access to the system. Each user has a user information profile to store information about them. You can assign users to user groups. When you assign a role directly to a user, all access rights defined in the role are inherited by the user.

Figure 37: User and resource management shows the user and resource management architecture.

Figure 37: User and resource management



## 13.2.1 Roles

Roles define the application access and resource permissions that can be assigned. You first create roles, then associate them to user groups according to the type of network activities the user group is meant to perform. Each member user of a user group can perform the roles specified for that group.

Optionally, you can also choose to assign a role directly to a user. When a user requires a specific set of permissions, you can bypass the use of user groups entirely.

Each role is mapped to a specific set of resource group access permissions. When a role is created, you can set the possible levels of permission for the associated resource groups to any the following:

<b>No Access</b>	Users or user group members do not have access to this resource group. The No Access permission is set by default for each resource until you change it when defining the role.
<b>Read</b>	Read permissions allow users or user group members to view specific resources, but they cannot make changes.
<b>Read / Write</b>	Read/write permissions allow users or user group members to view and modify resources.

After a role is created, you can return to the role and modify the resource access permissions.

### 13.2.1.1 Viewing a list of existing roles

#### About this task

Follow this procedure to view a list of existing roles.

#### Procedure

**Step 1.** Click  to open the main menu.

**Step 2.** Select  **User and Resource Management**.

**Step 3.** Select **Roles** from the drop-down list.

### 13.2.1.2 Predefined roles

The following table shows the predefined system roles and describes the specific permissions each role allows users. These roles are defined with common resource access privileges that you can quickly assign to new users. Administrators can associate roles to a specific user or to all members of user groups.

Predefined roles cannot be modified. To create customized roles for users that require specific permissions, see [Creating a role](#).

Table 23: Predefined roles

Role	Description
fabric-operator	Allows read/write access to all system resources except infracomponents (such as users, roles, and resource groups).
fabric-viewer	Allows read only access to system resources.
fss-admin	Allows admin privileges for all system resources in default namespaces.
ztp	Allows access to node management resources.

### 13.2.1.3 Creating a role

#### About this task

Follow this procedure to create a new role.

#### Procedure

**Step 1.** Open the **Roles** list as described in [Viewing a list of existing roles](#).

**Step 2.** Click **+ CREATE ROLE**.

**Step 3.** Under the Role Info heading, specify a role name and add an optional description to describe the purpose of the role.

**Step 4.** Specify the resource access permissions for the role. For a specific resource group permission, select one of the following options from the drop-down list.

- Read
- Read / Write

Do this for one or more resource groups.

The No Access permission is automatically selected for each resource until you change it.

You can also use the sort and filter columns to narrow the resource access options, as described in [Lists](#).

**Step 5.** Click **CREATE**.

### 13.2.1.4 Modifying the resource access permissions of a role

#### About this task

After a role is created, you can modify its resource access permissions. You can specify the resources in the system that can be accessed by users or user group members with the role applied.

Follow this procedure to modify the application access permissions of a role.

#### Procedure

**Step 1.** Open the **Roles** list as described in [Viewing a list of existing roles](#).

**Step 2.** For a specific role, click the  options menu at the right end of the row.

**Step 3.** Select **Open**.

**Step 4.** Under the Resource Access heading, for a specific resource group permission, select one of the following options from the drop-down list.

- No Access
- Read
- Read / Write

Do this for one or more resource groups.

You can also use the sort and filter columns to narrow the resource access options, as described in [Lists](#).

**Step 5.** Click **SAVE**.

### 13.2.1.5 Deleting a role

#### About this task

Follow this procedure to delete a role.

#### Procedure

**Step 1.** Open the Roles menu as described in [Viewing a list of existing roles](#).

**Step 2.** For a specific role, click the  options menu at the right end of the row.

**Step 3.** Click **Delete...**

If prompted, confirm that you want to delete the selected role.

## 13.2.2 Resource groups

Resource groups define the specific system resources that users can access in the UI. Resource groups are associated with roles. When administrators assign a role to a user group, the user group gains access to the network resources specified in the resource group. Resource groups represent the ability to configure specific types of things in the system, such as fabrics, workloads, or profiles.

Resources are organized into groups based on logical functional boundaries. Access to the specific functional areas can be associated with roles, and ultimately to user groups according to the type of network activities the user group is meant to perform.

The system comes with a set of predefined resource groups for common functional areas. Administrators cannot create new resources groups.

### 13.2.2.1 Viewing a list of resource groups

#### About this task

Follow this procedure to view a list of existing resource groups.

#### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** Select  **User and Resource Management**.
- Step 3.** Select **Resource Groups** from the drop-down list.  
Use the sort and filter columns to narrow the list, as described in [Lists](#).

### 13.2.2.2 Predefined resource groups

[Table 24: Predefined resource groups](#) shows the predefined resource groups and describes the specific permissions that the resource groups allow for users. When creating a role, administrators can associate a specific resource group or a combination of multiple resource groups.

For example, the administrator may want to allow only a specific set of users to configure QoS policies. The administrator can create a role that includes the **RG-qos** resource group, with the permission set to **read / write**. Then, the administrator can create a user group with this particular role assigned. User members of the user group are allowed permission to view and modify QoS profiles.

*Table 24: Predefined resource groups*

Resource group	Resource type	Description
RG-AAA	AAA	User, user group, and resource group management
RG-alarmmgr	alarmmgr	Access to system alarm manager objects
RG-catalog	catalog	Access to catalogs
RG-connect	connect	Access to system connect objects
RG-devices	device	Configuring devices
RG-fabric	fabric	Configuring fabrics between the switches
RG-fabric-telemetry	fabric-telemetry	Access to all the statistics, events, and so on that are collected between the switches.
RG-fabric-uplinks	fabric-uplinks	Configuring the fabric uplinks of a fabric
RG-fabric-uplink-protocols	fabric-uplink-protocols	Configuring the uplinks of the fabric, which can include Layer-2, Layer-3, and protocols to interface with external routers connected to fabric uplinks
RG-images	image	Provides OS images of the device

Resource group	Resource type	Description
RG-infra	infra	Access to infra settings
RG-label	label	Access to labels
RG-Layer-1	layer1	Configuration of Layer-1 aspects such as SFP, breakout, and so on
RG-Layer-2	layer2	Layer-2 configurations such as VLANs, LAG, and so on
RG-Layer-2-protocols	l2protocols	Layer-2 protocol configurations such as LLDP, LACP, and so on
RG-Layer-3	layer3	Layer-3 configurations such as sub-interfaces, static and dynamic routing policies, and so on
RG-Layer-3-protocols	l3protocols	Layer-3 protocol configurations such as BGP, OSPF, ISIS, and so on
RG-Maintenance	maintenance	Access to node maintenance intent
RG-Management	management	Configuring the management VRF and the relevant CoPP aspects
RG-mgmt-protocols	management-protocols	Configuring protocols used to manage SR Linux such as SSH, gNMI, NTP, FTP, and so on
RG-qos	qos	Configuration of QoS policies; can include CoPP
RG-region	region	Creating and configuring data center regions
RG-sandbox	sandbox	Access to sandbox environment
RG-security	security	Configuring security policies for workloads and fabrics
RG-topology	topology	Topology access for fabrics
RG-workload	workload	Configures workload related policies, including ACL and QoS policy profiles, but cannot edit the ACL and QoS profiles
RG-workload-attachments	workload-attachments	Configure workload attachment points
RG-workload-telemetry	workload-telemetry	Access to all statistics, events, flows, and so on, collected on the downlinks

### 13.2.3 User groups

A user group associates multiple users with a role, enabling them to access network resources. Administrators can create user groups and assign a specific role to each group according to the type of network activities the user group is meant to perform. When a role is assigned to a user group, all users within the group have the same access to resources, as specified by the role.

You can assign multiple users to a group. Users can also be members of multiple groups.

### 13.2.3.1 Viewing a list of existing user groups

#### About this task

Follow this procedure to view a list of existing user groups.

#### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** Select  **User and Resource Management**.
- Step 3.** Select **User Groups** from the drop-down list.

### 13.2.3.2 Predefined user groups

The following table shows the predefined system user groups. Each member of the user group is assigned the corresponding role and gains the resource access privileges associated with the role. These user groups represent collections of users that have been assigned specific system predefined roles. See [Predefined roles](#).

To create customized user groups for users that require specific permissions, see [Creating a user group](#).

Table 25: Predefined user groups

User Group Name	Role	Description
fabric_operator_group	fabric-operator	Allows group members read/write access to all system resources except infra components (such as users, roles, and resource groups).
fabric_viewer_group	fabric-viewer	Allows group members read only access to system resources.
fss_admin_group	fss-admin	Allows group members admin privileges for all system resources in default name spaces.
ztp_group	ztp	Allows group members access to node management resources.

### 13.2.3.3 Creating a user group

#### About this task

A user group can be configured initially without assigning users. If you want to add users to the group, you should also have some users configured. See [Users](#).

Follow this procedure to create a new user group.

## Procedure

- Step 1.** Open the **User Groups** menu as described in [Viewing a list of existing user groups](#).
- Step 2.** Click **+ CREATE USER GROUP**.
- Step 3.** Under the **User Info** heading, specify a user group name and add an optional description to describe the purpose of the group.
- Step 4.** Under the **Assigned Users** heading, select users to add to the user group. Click the check box next to a user's name to add them to the group.  
Optionally, you can create a user group without users, then add users later.  
You can sort and filter the list of users as described in [Lists](#).
- Step 5.** Click **CREATE**.

### 13.2.3.4 Assigning the specified role of a user group

#### About this task

You can assign a role to a user group by editing an existing user group. A role can also be assigned when you create a group.

You can only assign one role per user group. If a user group has an assigned role, you can change the assignment.

Follow this procedure to assign a role to a user group.

#### Procedure

- Step 1.** Open the **User Groups** menu as described in [Viewing a list of existing user groups](#).
- Step 2.** For a specific user group, click the  options menu at the right end of the row.
- Step 3.** Select **Open**.
- Step 4.** Select **Assigned User Roles** from the left navigation menu.  
You can sort and filter the list of user roles as described in [Lists](#).
- Step 5.** Select one of the listed roles to assign to the user group.  
You can only select one role.
- Step 6.** Click **SAVE**.

### 13.2.3.5 Deleting a user group

#### About this task

Follow this procedure to delete a user group.

#### Procedure

- Step 1.** Open the Users Groups menu as described in [Viewing a list of existing user groups](#).
- Step 2.** For a specific user group, click the  options menu at the right end of the row.
- Step 3.** Click **Delete...**  
If prompted, confirm that you want to delete the selected user group.

## 13.2.4 Users

Users are individuals with access to the system. Users gain access to application and network resources through the user groups to which they are assigned.

Individual users can also be assigned roles directly, without membership to a user group.

### 13.2.4.1 Viewing a list of existing users

#### About this task

Follow this procedure to view a list of existing users.

#### Procedure

- Step 1.** Click  to open the main menu.
- Step 2.** Select  **User and Resource Management**.
- Step 3.** Select **Users** from the drop-down list.

### 13.2.4.2 Creating a new user

#### About this task

Follow this procedure to create a new user.

#### Procedure

- Step 1.** Open the Users list as described in [Viewing a list of existing users](#).
- Step 2.** Click **+ CREATE USER**.
- Step 3.** In the User Info section of the overlay, enter the required job information data for the new user.
  - a.** a user name to identify the user
  - b.** the user's e-mail address
  - c.** a password the user uses to sign in to the system
  - d.** confirm the password
  - e.** optionally, enter the user's first and last name to identify them
- Step 4.** Confirm the **User Enabled** toggle is enabled.

You can also disable the user by disabling the toggle. Only enabled users are permitted access to the system.
- Step 5.** Assign the user to one or more user groups. Select **Assigned User Groups** from the left navigation menu.

The user must be assigned to at least one user group to gain resource permissions.

Optionally, you can create a user without assigning the user to a user group, then add the user to a user group later.
- Step 6.** Click the check box next to one or more user groups to which you would like to assign the user.
- Step 7.** Click **CREATE**.

### 13.2.4.3 Assigning a user to a user group

#### About this task

You can assign a user to one or more user groups. The user gains the resource permissions that are associated with the group. You should configure appropriate user groups before starting this procedure. Follow this procedure to assign a user to a group.

#### Procedure

- Step 1.** Open the Users list as described in [Viewing a list of existing users](#).
- Step 2.** For a specific user, click the  options menu at the right end of the row.
- Step 3.** Select **Open**.
- Step 4.** Select **Assigned User Groups** from the left navigation menu.  
You can sort and filter the list of user groups as described in [Lists](#).
- Step 5.** Click the check box next to one or more user groups to which you would like to assign the user.
- Step 6.** Click **SAVE**.

### 13.2.4.4 Assigning a role to a user

#### About this task

You can assign a role directly to a user, bypassing user groups entirely. The user gains the resource permissions specified in the role. Only one role can be assigned to an individual user. If the user requires customized permissions that are not present in any of the existing roles in the system, you can create a custom role for the user containing the specific permissions. See [Creating a role](#).

Follow this procedure to assign a role directly to a user.

#### Procedure

- Step 1.** Open the Users list as described in [Viewing a list of existing users](#) and select a user.  
If the user that you would like to assign the role to is not yet created in the system, add the user as described in [Creating a new user](#).
- Step 2.** For the specific user, click the  options menu at the right end of the row.
- Step 3.** Select **Open**.
- Step 4.** Select **Assigned User Roles** from the left navigation menu.  
You can sort and filter the list of user roles as described in [Lists](#).
- Step 5.** Select one of the listed roles to assign to the user.  
You can only select one role for a user.
- Step 6.** Click **SAVE**.

#### What to do next

To confirm the assignment, you can open the role and view the list of assigned users.

### 13.2.4.5 Deleting a user

#### About this task

Follow this procedure to delete a user from the list of users.

#### Procedure

**Step 1.** Open the Users list as described in [Viewing a list of existing users](#).

**Step 2.** For a specific user, click the  options menu at the right end of the row.

**Step 3.** Click **Delete...**

If prompted, confirm that you want to delete the selected user.

#### Expected outcome

The user is deleted from the list of users. The user is also removed from any user group memberships.

## 13.3 Backup and restore

The Fabric Services System uses Kubernetes' ability to take a point-in-time (PiT) snapshot to implement backup and restore capabilities. A snapshot represents a point-in-time copy of a volume. A snapshot can be used to provision a new volume or to restore the existing volume to the previous state captured in the snapshot.

To take these point-in-time snapshots, the container storage interface (CSI) standard is used. This standard allows the creation and deletion of volume snapshots via the Kubernetes API and the creation of new volumes pre-populated with the data from a snapshot via Kubernetes dynamic volume provisioning.

The Fabric Services System uses the BorgBackup program to move the point-in-time snapshot to a remote location. For more information about the BorgBackup program, see <https://borgbackup.readthedocs.io/en/stable/index.html#>. Through BorgBackup, the backup solution supports:

- Data de-duplication to lower the amount of storage used
- Data encryption to secure the backup
- (Optional) Remote storage of the backup over SSH on a remote server (using SSHFS)

#### Backup consistency in a micro-service architecture

A backup must be consistent across all the micro-services of an application to guarantee the restore capabilities. To achieve consistency, before the snapshots can be taken, ensure that all the micro-services have finished their write activity and that no further changes are accepted or incoming to the system. The best way to achieve this is by blocking access to the services temporarily.

The backup script has an option (**-s**) to specifically guarantee this consistency by scaling down all the micro-services for a short duration (the time to take the snapshots). Using this option prevents any changes to be made while the snapshot is taken. Nokia strongly recommends the use of this option; it is a requirement in production environments. When you omit this option, the restore functionality cannot be guaranteed.

Additionally, execute a backup only when:

- No background tasks are running like configuration generation or deployment of a fabric or workload intent.

- No Digital Sandbox Fabrics are present and no Digital Sandbox workloads are running.

The deployer VM must also be available during the backup procedure to ensure that the Kubernetes cluster has access to the registry with the image of all services, in case they are needed.

### Supported restore scenarios

The restore operation from a backup always requires a fresh installation of the entire Fabric Services System deployment. The restore takes place as part of the installation of a fresh environment. The restore process has the following requirements:

- The same version of the Fabric Services System must be used; for example, if a backup is taken from a deployment with version 22.4.1-90, the restore has to use that same version.
- The Fabric Services System nodes on which the restore is executed must have the same IP addresses and FQDNs as the environment from which the backup was taken.
- The `input.json` file that was used for the installation of the environment from which the backup was taken must be used for the restore procedure.

### Backup and restore scripts

The backup and restore commands are executed from one of the Kubernetes nodes, preferably one of the master nodes. The node on which the command is executed must have passwordless (key-based) SSH access to all other Kubernetes nodes as the root user. If a remote server is used through the `FSS_BACKUP_REPOS` environment variable, passwordless (key-based) SSH is also required to the remote server.

You use the **fss-backup.py** script for the backup operation.

```
fss-backup.py [-h] -b backup_name [-s] [-i] [-l [location]] [-c]
```

You use the **fss-restore.py** script for the restore operation.

```
fss-restore.py [-h] -b backup_name [-l [location]] [-c]
```

Table 26: Parameter descriptions

Option	Specifies
<b>-h</b>	Display the help screen and describes the usage of the script and CLI arguments.
<b>-b backup_name</b>	Name of the backup file, consisting of lower-case alphanumeric characters, '-' or '.', and must start and end with an alphanumeric character.
<b>-s</b>	Enable automatic scale down of the Fabric Services System deployment before taking the snapshots. This option guarantees that all services are temporarily down and a consistent backup of the entire environment is taken. This option is strongly recommended.
<b>-i</b>	Include the logs volume in the backup.
<b>-l location</b>	Full path to the folder where the backup file is stored, for example <code>/data/backups</code> . This option overwrites and takes precedence over the <code>FSS_BACKUP_REPOS</code> environment variable. This folder is local on the Kubernetes node where the backup command is executed, but can be mounted from a remote location like an NFS storage. Use the following format:

Option	Specifies
	<ul style="list-style-type: none"> <li>if the backup file is stored in a local directory: <code>/&lt;full path&gt;/&lt;directory&gt;</code>, for example, <code>/data/backups</code></li> <li>if the backup file is stored in a remote location: <code>user@&lt;ip address hostname&gt;:/&lt;full path&gt;/&lt;directory&gt;</code></li> </ul>
<b>-c</b>	Clean up a failed backup.

The scripts read the following environment variables:

- BORG\_PASSPHRASE** - a required environment variable; this password string is required by the Borg Backup utility to encrypt or decrypt backup files. Keep this passphrase safe and secure, if you lose it, access to the backup data is lost as well.
- FSS\_BACKUP\_REPOS** - specifies the location where the backup files are stored. It can be a local directory or remote location. You can use the **-l** option to override the setting for the **FSS\_BACKUP\_REPOS** variable.
- BACKUP\_RESTORE\_LOGS** - specifies log location for the backup and restore operation. If it is not configured, logs are stored in logs directory where scripts are being executed.

In the following example, compute6 is used as the backup repository.

```
export BORG_PASSPHRASE='XYZ2022_n0kIaNeTw0rKs_FsS1!&&123'
export FSS_BACKUP_REPOS=root@compute6:backup-repo
export BACKUP_RESTORE_LOGS=/root/backup-restore/logs
```

### 13.3.1 Backing up

#### About this task

The resulting backup is a full backup.

#### Procedure

**Step 1.** Run the backup script on the master node of the Kubernetes cluster that needs to be backed up.

```
/root/backup-restore/fss-backup.py -b <backup_name> [-s] [-i] [-l <location>]
```

#### Example

```
/root/backup-restore/fss-backup.py -b 0427-fss1 -s -l /root/backups
```

**Step 2.** Display the status of the backup operation.

#### Example

You can view the status of the backup in the `status.json` file in the logs directory.

```
[root@master-node backup-restore]# cat logs/0427-fss1-logs/status.json
{
  "fss_version": "v22.4.1-52",
  "status": "backup_completed",
  "completed_pvcs": [
```

```

"data-dev-postgresql-0",
"datadir-0-dev-cp-kafka-0",
"datadir-0-dev-cp-kafka-1",
"datadir-0-dev-cp-kafka-2",
"datadir-dev-cp-zookeeper-0",
"datadir-dev-cp-zookeeper-1",
"datadir-dev-cp-zookeeper-2",
"datadir-dev-neo4j-core-0",
"datalogdir-dev-cp-zookeeper-0",
"datalogdir-dev-cp-zookeeper-1",
"datalogdir-dev-cp-zookeeper-2",
"dev-fss-dhcpd",
"dev-fss-dhcpd-lease",
"dev-fss-image-mgmtstack",
"dev-fss-node-checkpoints",
"dev-fss-nodecfg",
"dev-mongodb"
],
"failed_pvcs": [],
"failure_msg": "",
"last_updated_at": "2022-04-28T04:09:30.453392+00:00"
}

```

**Step 3.** Verify the backup files in the repository.

#### Example

```

[root@master-node backups]# borg list 0427-fss1 --short
manifests
data-dev-postgresql-0
datadir-0-dev-cp-kafka-0
datadir-0-dev-cp-kafka-1
datadir-0-dev-cp-kafka-2
datadir-dev-cp-zookeeper-0
datadir-dev-cp-zookeeper-1
datadir-dev-cp-zookeeper-2
datadir-dev-neo4j-core-0
datalogdir-dev-cp-zookeeper-0
datalogdir-dev-cp-zookeeper-1
datalogdir-dev-cp-zookeeper-2
dev-fss-dhcpd
dev-fss-dhcpd-lease
dev-fss-image-mgmtstack
dev-fss-node-checkpoints
dev-fss-nodecfg
dev-mongodb
[root@master-node backups]#

```

### 13.3.2 Restoring a backup and installing the Fabric Services System application

The restore process consists of the following high-level tasks:

1. Set up new Fabric Services System nodes with the same IPs and FQDNs.
2. Install the Kubernetes cluster only (not the application) using the same `input.json` that was used for the original environment.
3. Restore the backup data using one of the Kubernetes nodes (preferably a master node).

4. Install the Fabric Services System application components using the same `input.json` that was used for the original environment.

### 13.3.2.1 Setting up new Fabric Services System nodes

#### About this task

The restore of the Fabric Services System backup has to be executed in a clean environment that uses the same IPs and FQDNs of the original environment. You can use the procedures from the *Fabric Services System Software Installation Guide* to set up such an environment.

#### Procedure

- Step 1.** Optional: Complete the procedure "Deploying and configuring the Fabric Services System deployer VM".



**Note:** Complete this procedure only if the deployment VM was also lost or needs to be reinstalled. Ensure that you are using the appropriate version, that is, the same version that was used to install original environment.

- Step 2.** Complete the procedure for you deployment scenario.

- "Virtual machine-based Installation: Using the Fabric Services System base OS image"
- "Bare metal-based installation: Preparing the Fabric Services System nodes"

#### Expected outcome



**Notice:** When the nodes have been installed, do not execute "Installing Fabric Services System" in the *Fabric Services System Software Installation Guide*.

### 13.3.2.2 Installing the Kubernetes cluster

#### About this task

Execute this procedure on the deployer VM using the same `input.json` that you used for the installation of the original deployment.

#### Procedure

- Step 1.** Initiate the setup.

#### Example

```
[root@fss-deployer ~]$ /root/bin/fss-install.sh configure sample-input.json
```



**Notice:** Do not execute the `fss-install.sh` script as suggested by the output of the preceding command.

- Step 2.** Start the installation of Kubernetes.

**Example**

```
[root@fss-deployer ~]$ /root/bin/setup-k8s.sh
```

The installation time varies depending on the capacity of your system.

**13.3.2.3 Restoring a backup****Prerequisites**

- Ensure that the backup folder created by the backup script is available to the restore command on the Kubernetes node on which you are executing this procedure. You can configure this through the FSS\_BACKUP\_REPOS environment variable pointing to a remote server using SSH/SSHFS or by putting the backup folder on the Kubernetes node itself and using the `-l` option, similar to how the backup was taken.
- Ensure that the BORG\_PASSPHRASE environment variable is set to the correct passphrase, otherwise, the restore will fail.

**Procedure**

**Step 1.** Restore the backup on the target Kubernetes node.

```
/root/backup-restore/fss-restore.py -b backup_name [-l location]
```

**Example**

```
./root/backup-restore/fss-restore.py -b 0427-fss1 -l /root/backups
```

The script does following:

- a. Checks if the specified backup file name exists in the specified location in the Borg repository; if the file is not present, the script terminates
- b. Mounts the Borg repository locally to read the backup content
- c. Creates PVCs from the YAML files in the manifest archive in the Borg repository that contains the PVC definition for all the original PVCs; also creates PVs of the same size as the original PVs in the source cluster
- d. Creates a mount pod for each PVC on the node where the restore script is being executed, then mounts newly created empty PVCs
- e. Copies the volume content of the mounted backup to the corresponding PVC volume
- f. Deletes the PVC mount pod; the PVCs and PVs are still present and point to the new data volume where data is copied

**Step 2.** Check the status of the restore operation in the `status.json` file in the `logs/<backup-name>-restore-logs` directory.

**Example**

```
[root@compute1 backup-restore]# cat logs/0501-3-fssbackup-restore-logs/status.json
{
  "fss_version": "v22.4.1-52",
  "status": "restore_completed",
```

```

"restored_pvcs": [
  "datadir-mvsrlfsp05-cp-zookeeper-2",
  "datadir-mvsrlfsp05-cp-zookeeper-0",
  "mvsrlfsp05-fss-nodectfg",
  "datadir-mvsrlfsp05-mongodb-secondary-0",
  "datalogdir-mvsrlfsp05-cp-zookeeper-0",
  "mvsrlfsp05-fss-image-mgmtstack",
  "mvsrlfsp05-fss-dhcpd-lease",
  "datalogdir-mvsrlfsp05-cp-zookeeper-1",
  "datadir-mvsrlfsp05-neo4j-core-0",
  "mvsrlfsp05-fss-node-checkpoints",
  "datalogdir-mvsrlfsp05-cp-zookeeper-2",
  "mvsrlfsp05-fss-dhcpd",
  "datadir-mvsrlfsp05-cp-zookeeper-1",
  "datadir-mvsrlfsp05-mongodb-primary-0",
  "mvsrlfsp05-fss-logs-volume",
  "data-mvsrlfsp05-postgresql-0",
  "datadir-0-mvsrlfsp05-cp-kafka-1",
  "datadir-0-mvsrlfsp05-cp-kafka-2",
  "datadir-0-mvsrlfsp05-cp-kafka-0"
],
"failure_msg": "",
"last_updated_at": "2022-05-01T19:10:26.207653+00:00"
}[root@mvsrlfsp05-compute1 backup-restore]#

```

**Step 3.** Verify the restored files.

Execute the **kubectl get pvc** command. The output should match the output of the **borg list** command in step 3 of [Backing up](#).

#### Expected outcome

When the output shows that the PVC have been installed, you can install the Fabric Services System application.

### 13.3.2.4 Installing the Fabric Services System application

#### About this task

Now that the backup data has been restored, you can install the Fabric Services System in the Kubernetes cluster. Use the deployer VM from which you executed the commands to install the Kubernetes cluster.

#### Procedure

**Step 1.** Initiate the setup.

#### Example

```
[root@fss-deployer ~]$ /root/bin/fss-install.sh configure sample-input.json
```



**Notice:** Do not execute the **fss-install.sh** script as suggested by the output of the preceding command.

**Step 2.** Start the installation of the Fabric Services System application.

**Example**

The installation time varies depending on the capacity of your system.

```
[root@fss-deployer ~] $ /root/bin/fss-app-install.sh
```

**Expected outcome**

At the end of this procedure, the Fabric Services System application has been restored from your backup and is ready for use.

## 13.4 Capturing troubleshooting data

If the Fabric Services System application suffers any failure you would like to report to Nokia technical support, you can use a script that is bundled with the system to generate information about the Fabric Services System application status. You can then share this information with Nokia support teams to assist in troubleshooting the problem.

The technical support script is included with the Fabric Services System, provided you deployed the system using the standard deployer Virtual Machine (VM).

You can execute the script from the installer VM, or from any Kubernetes (K8s) master node (that is, the compute node that is marked as master explicitly, or the first node in the input configuration file). For details about the installer VM and the nodes participating in the Fabric Services System cluster, see the *Fabric Services System Software Installation Guide*.

The script creates a file containing status information about the Fabric Services System microservices and CPU statistics that can be useful when troubleshooting issues.

### 13.4.1 How to run the technical support script

The technical support script can be found on the installer VM or any K8s master node, and is named one of the following:

- on the installer VM: `deployer-tech-support.sh`  
Running the script on the installer VM is particularly useful if an error arises during installation of the Fabric Services System.
- on the K8s master node: `techsupport.sh`

For the current release, the script supports the following options:

**help**

displays a list of supported option and their descriptions

**minimal**

generates a `.tar` file containing basic system status output.

**complete**

generates a `.tar` file containing comprehensive system status output.

**db mongo**

generates a folder containing a backup of the mongo database.

**filter-criteria**

provides a sub-string to select pod names (for example, entering "fss-da" causes the script to gather data only for those pods whose name includes "fss-da").

**encryption**

encrypts the file so that the information is secure until accessed by Nokia technical support using its key.

For example, when running the script from the k8s master node:

```
bash techsupport.sh minimal
```

```
bash techsupport.sh db mongo
```

The script also creates a folder named debug-dump, into which it places each of the files generated by the data collection process.

When the script is complete, it generates a .tar file in the debug-dump folder whose name is the current date and time. For example: Mon-Jul-26-18-08-31-IST-2021.tar.gz.

You can optionally encrypt the file so that only the Nokia technical support team can unlock it with their key. For example:

```
bash tech-support.sh encrypt Mon-Jul-26-18-08-31-IST- 2021.tar.gz
```

This generates a file named Mon-Jul-26-18-08-31-IST- 2021.tar.gz.enc, which you can then share with the Nokia support team.

## 13.5 Node discovery

For the Fabric Services System to serve its function of building and managing a data-center fabric, it must be able to discover, communicate with, and configure any SR Linux nodes that are intended to be members of that fabric. The system accomplishes this using a variation of the standard SR Linux Zero-touch Provisioning (ZTP) process.

As described in [Fabric intents](#), node discovery is a key prerequisite before the system can deploy a fabric intent. Each node you include within a fabric intent must first be discovered and have reported itself to be in a Ready state before the system can deploy the additional configuration details to make the node adopt its assigned role in the fabric. This appendix describes how node discovery occurs, and how the node's Ready state is achieved.

This section summarizes the SR Linux ZTP process and highlights the elements that are new for the Fabric Services System. For additional details about the SR Linux ZTP process, see the *Nokia Service Router Linux Software Installation Guide*.

### 13.5.1 Participants

Several entities participate in the ZTP process when used in conjunction with the Fabric Services System.

- The Fabric Services System: this process assumes that the system is up and running and able to communicate across the network. Based on a user's input when designing a fabric intent, the system creates a set of configuration files that are used to configure each node so that it can perform its intended role in the fabric.

Two subsystems integrated within the Fabric Services System also participate in the ZTP process:

- a DHCP server: the system's integrated DHCP server is capable of receiving a newly booted SR Linux node's DHCP request and replying with an assigned IPv4 address, provided this SR Linux serial number is assigned to a fabric.  
You can use an external DHCP server if you prefer.
- an HTTP file server: the system's integrated HTTP file server stages and transfers SR Linux image files, scripts, and other node configuration files after communication with an SR Linux node is established.  
You can use an external HTTP file server if you prefer.
- The SR Linux node: when first booted up, an SR Linux node continuously sends DHCP discovery messages requesting an IP address. After it receives an address and subsequent data from a DHCP server, ZTP software built into the Fabric Services System can process this information and allows the node to configure itself in accordance with information sent by a management system (in this case, the Fabric Services System).

### 13.5.2 The discovery and configuration process

The following sequence describes the series of events from the initial fabric intent design, to the communication with a newly installed SR Linux node, to the configuration of that node so it can function in its assigned role in the planned fabric.

Steps 2 through 7 in this process represent the standard SR Linux ZTP process.

Step 8 configures support for ZTP with the Fabric Services System. This step specifies the system server, its role as the fabric manager, and the FTP server's IP address and authentication data. These settings allows the node to communicate directly with the system and to be configured as required to participate in the planned fabric design.

1. In the Fabric Services System, an operator:
  - a. Adds any required SR Linux image intended for use on a fabric's nodes to the file server.
  - b. Creates a fabric, specifying an SR Linux image.
  - c. Associates individual nodes in the fabric intent with planned hardware, based on the hardware serial number.

The system waits for the required nodes to communicate with the DHCP server.

2. A newly-booted SR Linux node comes online and sends out IPv4 DHCP discovery requests.
3. The system's internal DHCP server receives the request, replies with the assigned IP address, and points the SR Linux node to a ZTP configuration file (*provisioning.py*) on the integrated file server.
4. The node requests the ZTP configuration file from the system's file server and receives the file in reply. The configuration file identifies a particular SR Linux image for the node to use, and an associated MD5 file.
5. The node requests the MD5 file from the system's file server and receives the file in reply.
6. The node requests the binary file containing the correct SR Linux image from the file server, and receives the file in reply.
7. The node compares the received SR Linux image version to its current SR Linux version. If necessary, the node upgrades or downgrades itself to the new version of SR Linux contained in the image file.

8. The node requests from the file server the `initconfig.json` file identified in the `provisioning.py` file, and receives the file in reply. This file identifies the Fabric Services System as the fabric manager and includes the IP address and authentication information that the node can use to communicate with the system.
9. The node "phones home," telling the Fabric Services System server that the node is configured and ready for management.
10. Internally, the system updates the node's status to Ready.
11. After all nodes are in the Ready state, the operator must:
  - a. Add the fabric to region's deployment pipeline.
  - b. Deploy the fabric from the deployment pipeline.

The system deploys additional configuration data to the node, so that the node can configure itself for participation in the fabric as intended by the fabric design from step 1.

### 13.5.3 Using an external DHCP server

#### About this task

You can use an external DHCP server instead of the one integrated into the Fabric Services System to discover and configure the nodes to participate in a fabric intent.



**Note:** When using an external DHCP server, set the region's **Use internal DHCP** property to Disabled.

In this case, the process to configure an SR Linux node follows the usual Zero Touch Provisioning process documented in the *Nokia Service Router Linux Software Installation Guide*. However, the configuration that your DHCP server loads onto the new nodes must be the initial configuration that was generated by the Fabric Services System when you designed the fabric intent that contains it.

You use these node configurations as part of the node discovery and configuration process conducted using your DHCP server and otherwise following the Zero Touch Provisioning process documented in the *Nokia Service Router Linux Software Installation Guide*.

You also need to configure the nodes so that when they boot up, they "call home" to the Fabric Services System for subsequent management. For guidance on configuring nodes in this way when using an external DHCP server, please contact Nokia technical support.

To obtain the initial configuration for all of the nodes participating in a fabric intent, do the following:

#### Procedure

- Step 1.** Open the fabric intent as described in [Viewing a fabric intent](#).
- Step 2.** From the opened fabric intent, click the  menu at the upper right of the page.
- Step 3.** Select **Download Initial Node Configuration** from the list.

#### Expected outcome

The system immediately downloads the file "initialNodeConfigs" to your Downloads folder.

- Step 4.** To view the configuration, open the initialNodeConfigs file in a text editor.

## Appendix: Fabric intent parameters

This appendix describes the fabric intent parameters in the Fabric Services System. When creating a new fabric intent, you can modify these parameters to fine-tune the fabric intent design. You must initially enter some required parameters, then you can modify optional parameters to further customize the fabric intent.

### Required parameters

[Table 27: Required parameters \(fabric intent\)](#) defines the required fabric intent parameters and the available values that can be selected.

Some parameters apply only to some fabric intent types:

- L/S: "Y" indicates the parameter applies to flexible leaf/spine fabric intents.
- BB: "Y" indicates the parameter applies to Backbone fabric intents.

*Table 27: Required parameters (fabric intent)*

Parameter	L/S	BB	Description	Values/Range
Reference Template	Y	Y	Predefined Nokia certified templates to create a fabric intent.	Flexible Leaf Spine Backbone
Fabric Type	Y	Y	The type of networking switch fabric to be implemented.	Real Digital Sandbox
Fabric Intent Name	Y	Y	The name of the fabric.	A string value
Description	Y	Y	A description of the fabric intent.	A string value
Prefix Naming	Y	Y	A string to be added at the beginning of the name of every node in the fabric intent. The rest of the node name is automatically generated.  For example, enter "A01" here to assign nodes names such as "A01-leaf-1" and "A01-spine-1".  The Prefix Naming string must be unique to each fabric.	A string value
Region	Y	Y	Identifies the region that contains the new fabric intent. The Fabric Services System currently supports only a single region.	N/A

Parameter	L/S	BB	Description	Values/Range
			To create or edit a region now, you can click on the Regions icon (  ) to navigate to the Regions page. This exits the fabric intent creation procedure. You must begin fabric intent creation again after creating the new region.	
Software Catalog / Image ID	Y	Y	The version of SR Linux software to load onto all of the nodes in the new fabric.	Only software that has been loaded into the catalog. Currently supported: <ul style="list-style-type: none"> <li>SR Linux 21.11.2-72</li> <li>SR Linux 21.6.4-6</li> </ul>
Is Border Leaf	-	Y	Enable if this backbone fabric is serving as a "border leaf", able to carry EVPN workloads as well as provide access to an eternal network.  Disable if this is a standard backbone that is not capable of carrying an EVPN workload.	Enabled, Disabled  If enabled, only the 7220 IXR-D3 is supported as a node type.
Spine Node Type	Y	-	Select a supported type of hardware to act as a spine node.	7220 IXR-D3 7220 IXR-D3L 7250 IXR-6 7250 IXR-10 7220 IXR-H2 7220 IXR-H3
Number of Spine Nodes	Y	-	Specify the number of devices of the specified type to act as spine nodes.	An integer
Add Pod	Y	-	Indicates whether you are configuring a Pod, which is a subset of the fabric that can use a different hardware type for leaf nodes than the non-Pod subset. For more information about Pods, see <a href="#">Adding a second type of hardware for leaf nodes</a> .	Enabled, Disabled  If Enabled, the configuration you create is applied to the Pod subset of your fabric. If Disabled, the configuration is applied to the non-Pod subset.
Leaf Node Type	Y	-	Select a supported type of hardware to act as a leaf node.	7220 IXR-D2 7220 IXR-D2L 7220 IXR-D3

Parameter	L/S	BB	Description	Values/Range
				7220 IXR-D3L
Number of Leaf Nodes	Y	-	Specify the number of devices of the specified type to act as leaf nodes.	An integer
Node Type	-	Y	Select a supported type of hardware to act as a backbone node.  The available values vary based on the "Is Border Leaf" selection.	7220 IXR-D2 7220 IXR-D2L 7220 IXR-D3 7220 IXR-D3L 7250 IXR-6 7250 IXR-10 7220 IXR-H2 7220 IXR-H3
Number of Nodes	-	Y	Specify the number of devices of the specified type to act as backbone nodes.	An integer
Dual Homing	-	-	A topology is generated where devices can connect to the network through two independent access points, where one point is the primary connection and the other point is activated in event of a failure.	Enabled, Disabled
Servers per Rack	-	-	The number of servers per rack in the fabric intent.	1-40 servers per rack for all templates
Number of Racks	-	-	The number of racks in the fabric intent. The maximum number of racks depends on the reference template selected.	Leaf Spine: 1-8 racks Chassis Spine: 1-20 racks Backbone: 1-30 racks Super Backbone: 1-40 racks
Edge Link Speed	Y	-	The data transfer speed from the access point to data center or computers.	1G, 10G, 25G, 40G, 50G, 100G, 200G, 400G
T1 Layer Uplink Speed	Y	-	The data transfer speed from computers or data center to the access point.	1G, 10G, 25G, 40G, 50G, 100G, 200G, 400G
Deploy Trigger Percentage	Y	Y	The percentage of fabric deployment for a fabric design intent to create a topology.	Has a fixed, automatically assigned value of 0% or 100% based on fabric type (Real or Digital Sandbox).

Parameter	L/S	BB	Description	Values/Range
				This parameter cannot be edited.

## Optional parameters

After choosing the required parameters, you can modify optional parameters to further customize the fabric intent. [Table 28: Optional parameters \(fabric intent\)](#) defines the optional fabric intent parameters and the values that can be modified.

Table 28: Optional parameters (fabric intent)

Parameter	Description	Values/Range
Tier One Oversubscription Parameters		
Over Subscription Ratio	The ratio between the bandwidth of the downlink ports and the bandwidth of the uplink ports for the first layer of the fabric.	Ratio of numbers with a colon between the two numbers, for example 1:3
Tier Two Oversubscription Parameters		
Over Subscription Ratio	The ratio between the bandwidth of the downlink ports and the bandwidth of the uplink ports for the second layer of the fabric.	Ratio of numbers with a colon between the two numbers, for example 1:3
Mgmt Traffic QoS Profile Parameters		
Dscp All Mgmt	QoS priority set to all DSCP management packets.	BESTEFFORT, LOPRIORITY, MIPRIORITY, HIPRIORITY This field is currently read-only.
Tos All Mgmt	QoS priority set to all ToS management packets.	BESTEFFORT, PRIORITY This field is currently read-only.
Dscp gNMI	QoS priority set to all DSCP gNMI packets.	BESTEFFORT, LOPRIORITY, MIPRIORITY, HIPRIORITY This field is currently read-only.
Dscp Img Download	QoS priority set to DSCP image downloads either FTP or HTTP-based.	BESTEFFORT, LOPRIORITY, MIPRIORITY, HIPRIORITY This field is currently read-only.

Parameter	Description	Values/Range
Dscp S Flow	QoS priority set to DSCP sFlow packets.	BEST_EFFORT, LOPRIORITY, MIPRIORITY, HIPRIORITY This field is currently read-only.
Dscp IP Fix	QoS priority set to DSCP IP fix packets.	BESTEFFECT, LOPRIORITY, MIPRIORITY, HIPRIORITY This field is currently read-only.
Uplink Parameters		
Uplink Speed	The uplink speed for a leaf-spine fabric.	1G, 10G, 25G, 40G, 50G, 100G, 200G, 400G
Unicast Cp	The type of unicast routing protocol used in the network topology.	BGP, OSPF, ISIS This field is currently read-only.
Multicast Cp	The type of multicast routing protocol used in the network topology.	None, PIMSM, PIMSSM, PIMBDIR This field is currently read-only.
Global Configuration parameters		
Enable FTP	Enables the File Transfer Protocol on the nodes in this fabric.	Enabled, Disabled
NTP	Specifies one or more Network Time Protocol (NTP) servers. Adding one or more NTP server IP addresses enables the system NTP client on the node, and is used for system clock synchronization.	IP address of one or more NTP servers

## Appendix: Workload VPN intent parameters

This appendix describes the workload intent parameters in the Fabric Services System.

### Workload intent

[Table 29: Basic parameters \(workload intent\)](#) and [Table 30: Subnet parameters \(workload intent\)](#) define the required and optional workload intent parameters and the appropriate values that can be inputted in the platform.

*Table 29: Basic parameters (workload intent)*

Parameter	Description	Values/Range
Workload VPN Intent Name	A unique name you assign to the workload VPN intent.	Any string value
Description	A description you provide for the workload VPN intent.	Any string value
Fabric Intents	Use this field to select one or more fabrics whose resources should participate in the workload intent.	By Fabric
Labels	Although not enabled during workload intent creation, this field can be used later to apply labels to the workload intent itself.	Supported labels

*Table 30: Subnet parameters (workload intent)*

Parameter	Description	Values
Name	A name you assign to the subnet.	Any string value
Description	A description you provide for the subnet.	Any string value
Type	Select a supported subnet type from the drop-down list.	Bridged Routed
IP Anycast Gateway (V4/V6)	For bridged subnets, an IP gateway to act as an IRB interface.	Enter a valid IPv4 or IPv6 address with a required CIDR.
IPv4 Learn Unsolicited ARP Enabled	For IPv4 addresses within the subnet, setting this to True enables the learning of ARP entries out of any ARP packet arriving at the IRB sub-interface, regardless of whether there was an ARP-Request issued from the IRB.	True, False
IPv6 Learn Unsolicited ARP Enabled	For IPv6 addresses within the subnet, setting this to True enables the learning of Neighbor Discovery Request entries	True, False

Parameter	Description	Values
	out of any Neighbor Discovery Request packet arriving at the IRB sub-interface, regardless of whether there was an Neighbor Discovery Request issued from the IRB.	
ACL Profile	For bridged subnets, an Access Control list that will restrict the traffic permitted to cross the subnet.	Select an already-configured ACL profile
IP MTU	For bridged subnets, the maximum transmission unit allowed.	1500+

Table 31: Sub-interface parameters (workload intent)

Parameter	Description	Values
Description	A description you provide for the selected sub-interface.	Any string value
VLAN ID	The VLAN ID associated with the current sub-interface.	Specify an available VLAN ID (1+)
ACL Profile	An Access Control List that defines the traffic that is permitted (and by implication, excluded) on the sub-interface.	Select an already-defined ACL profile.
Workload specific ACL optimization	Setting this value to On causes the system to verify that the selected ACL profile's IP address range is contained within the set of Gateway IP addresses you enter for this sub-interface.	On, Off
IP MTU	The Maximum Transmission Unit for the sub-interface; this is the maximum size for an IP packet that will not be fragmented in the course of transmission.	1500+
Association parameters		
Subnet	The subnet with which this sub-interface is associated.	Select an existing subnet from the drop-down list.
Association Type	The method used to associate this sub-interface with its "parent" subnet.	Currently fixed as "Node and Interface"
Node ID	The node within the fabric on which the current sub-interface is located.	Select an existing leaf node within the fabric or fabrics associated with this workload VPN intent.
Interface Name	The specific interface on the selected node with which this sub-interface is associated. This can be a LAG.	Select an interface from the drop-down list.
IP Gateway (V4/V6)	IP address of the forwarding device.	Enter the IP address of the gateway device.

Parameter	Description	Values
QoS parameters		
QoS Classifier (IP V4)	Quality-of-Service classifier (DSCP value and forwarding class) for an IPv4 address.	Select an already-defined QoS profile from the drop-down list.
QoS Classifier (IP V6)	Quality-of-Service classifier (DSCP value and forwarding class) for an IPv6 address.	Select an already-defined QoS profile from the drop-down list.
QoS Rewrite Rules (IP V4)	Quality-of-Service rewrite rules (forwarding class and DSCP selection) for an IPv4 address.	Select an already-defined QoS profile from the drop-down list.
QoS Rewrite Rules (IP V6)	Quality-of-Service rewrite rules (forwarding class and DSCP selection) for an IPv6 address.	Select an already-defined QoS profile from the drop-down list.

## Appendix: Supported alarms

### Equipment alarms

Alarm ID:	1001
Alarm name:	Fan Tray Fault
Description:	This alarm is raised when the associated fan-tray is in down/empty/failed/degraded/low-power operational state. The system may have cooling issues.
Severity:	Major
Probable cause:	Equipment malfunction
Remedial action:	The failed fan unit should be replaced.

Alarm ID:	1003
Alarm name:	Power Supply Fault
Description:	The alarm is raised when the associated power supply is not operationally Up. The specified power supply can no longer supply power to the system.
Severity:	Critical
Probable cause:	Power problem
Remedial action:	Check the status of the power supply.

Alarm ID:	1004
Alarm name:	Chassis Fault
Description:	The alarm is raised when chassis is operationally down.
Severity:	Critical
Probable cause:	Equipment malfunction
Remedial action:	Chassis Down

Alarm ID:	1005
Alarm name:	CPM Fault
Description:	This alarm is generated when the control module is in an operationally down/empty/failed/degraded/low-power state.
Severity:	Critical
Probable cause:	Equipment malfunction
Remedial action:	Remove the card and reset it. If this does not clear the alarm then please contact your Nokia support representative for assistance.

Alarm ID:	1006
Alarm name:	SFM Fault
Description:	The alarm is raised when the associated SFM module is in operationally down/empty/failed/degraded/low-power state. Traffic could be impacted.
Severity:	Critical
Probable cause:	Equipment malfunction
Remedial action:	The active CPM is at risk of failing to initialize after node reboot because it cannot access the SFM. Contact Nokia customer support.

Alarm ID:	1007
Alarm name:	Line Card Fault
Description:	The alarm is raised when the specified line card is operationally down/empty/failed/degraded/low-power state. Traffic is no longer being transmitted from this line card.
Severity:	Major
Probable cause:	Equipment malfunction
Remedial action:	Ensure line card is operationally up. Linecard may need to be replaced.

### Communication alarms

Alarm ID:	4001
Alarm name:	LLDP Adjacency Down
Description:	The alarm is raised when the Operational State of an LLDP adjacency is down. This is because the operational state of local interface is in a down state.
Severity:	Major
Probable cause:	DTE DCE Interface Error
Remedial action:	The operational state of the interface must be up in order for the selected adjacency to be up.

Alarm ID:	4002
Alarm name:	Interface Down
Description:	The alarm is raised when the operational state of interface is down.
Severity:	Critical
Probable cause:	DTE DCE Interface Error
Remedial action:	The condition exists because the physical interface is down either because it is administratively disabled, faulty or a cabling fault has occurred. Ensure that the interface is administratively up. Check for a poor cable connection to the port or for a faulty cable/fiber. If neither appears to be the problem run diagnostics on the port to determine if it is faulty.

Alarm ID:	4003
Alarm name:	Subinterface Down
Description:	The alarm is raised when the operational state of subinterface is down.
Severity:	Critical
Probable cause:	DTE DCE Interface Error
Remedial action:	The condition exists because the subinterface is down either because it is administratively disabled, faulty or a cabling fault has occurred. Ensure that the subinterface is administratively up. Check for a poor cable connection to the port or for a faulty cable/fiber. If neither appears to be the problem run diagnostics on the port to determine if it is faulty.

### Operational alarms

Alarm ID:	5001
Alarm name:	GNMI Connection Fault
Description:	GNMI connection to the network element has been lost.
Severity:	Major
Probable cause:	DTE DCE Interface Error
Remedial action:	Check network connectivity to restore GNMI connection.



# Customer document and product support



## Customer documentation

[Customer documentation welcome page](#)



## Technical support

[Product support portal](#)



## Documentation feedback

[Customer documentation feedback](#)