Cisco Crosswork Optimization Engine 1.1 User Guide

Last Modified: 2020-04-28

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

CHAPTER 1 **Overview of Cisco Crosswork Optimization Engine** 1

Audience 1

Overview of Cisco Crosswork Optimization Engine 2

Cisco Crosswork Optimization Engine Architecture 2

Crosswork Optimization Engine APIs 3

Segment Routing Path Computation Element (SR-PCE) 3

Log In and Log Out 3

Crosswork Optimization Engine Home Page 4

Set, Sort and Filter Table Data 7

CHAPTER 2 **Get Started** 11

Segment Routing 11

Inventory Management Concepts 13

Before You Begin 14

High-Level Workflows 15

Workflow: Auto-Onboard Devices 15

Workflow: Manually Import Devices 17

CHAPTER 3 **Manage Inventory** 19

Device Management Overview 19

About Adding Devices 19

Prerequisites for Onboarding Devices 21

Sample Configuration for Devices in Cisco NSO 22

Reachability and Operational State 23

Manage Credential Profiles 25

Create Credential Profiles 26

Import Credential Profiles 27
Edit Credential Profiles 30
Delete Credential Profiles 30
Export Credential Profiles 30
Change a Device's Credential Profile 31
Change the Credential Profile for Multiple Devices 31
Manage Providers 32
Add Cisco SR-PCE Providers 34
Auto-Onboard Property Descriptions 36
Cisco SR-PCE Reachability Issues 36
Configure Redundant Cisco SR-PCEs 37
Path Computation Client (PCC) Support 41
Add Cisco NSO Providers 42
Import Providers 43
Get Provider Details 45
Edit Providers 46
Delete Providers 46
Export Providers 47
View Devices Assigned to a Provider 47
Manage Network Devices 48
Import Devices 49
Add Devices Through the UI 51
Get Device Details 54
Filter Devices by Tags 56
Edit Devices 56
Delete Devices 57
Export Devices 57
View Device Job History 58
Manage Tags 59
Create Tags 60
Import Tags 61
Apply or Remove Device Tags 62
Delete Tags 62
Export Tags 63

CHAPTER 4 **Visualize the Network** 65

Network Topology Map 65

Troubleshoot Network Topology Map 67

Device and Link Icons 68

Configure Geographical Map Settings 69

Change the Layout of a Logical Map 69

Change Display Settings for Links, Devices, and TE Tunnel Metrics 70

Create Custom Map Views 72

Manage Custom Map Views 73

Visualize Devices 74

Get More Information About Devices on the Map 74

Access the Device Console 77

Identify the Members of a Cluster 78

Visualize Links 78

Get More Information About Links 78

Show Bandwidth Utilization for Links on the Map 80

Define Color Thresholds for Link Bandwidth Utilization 81

CHAPTER 5 **Visualize SR Policies and RSVP-TE Tunnels** 83

SR Policy and RSVP-TE Tunnel Support 83

SR Policy and RSVP-TE Tunnel Configuration Sources 88

PCC-Initiated SR Policy Example 88

PCC-Initiated RSVP-TE Tunnel Example 89

SR Policies and RSVP-TE Tunnels Topology Map 89

Highlight a TE Tunnel on the Map 92

Show Participating Nodes and Links 93

Show IGP, Delay, and Traffic Engineering Metrics 93

SR Policies Table 94

RSVP-TE Tunnels Table 96

Visualize SR Policies and RSVP-TE Tunnels 98

Visualize TE Tunnels Example 98

Configure Affinity Mapping 103

Preview Disjoint SR Policies and RSVP-TE Tunnels 104

View TE Tunnels Belonging to a Disjoint Group 107
Create and Manage SR Policies 107
Create Explicit Path SR Policies 108
Create Dynamic Path SR Policies 110
Modify SR Policies 113
Delete SR Policies 114
Get More Information About an SR Policy 114
Create and Manage RSVP-TE Tunnels 117
Create Explicit Path RSVP-TE Tunnels 117
Create Dynamic Path RSVP-TE Tunnels 118
Modify RSVP-TE Tunnels 120
Delete RSVP-TE Tunnels 121
Get More Information About an RSVP-TE Tunnel 121

CHAPTER 6 **Perform Administrative Tasks** 125
Manage Cisco Crosswork Network Automation 125
Monitor Cisco Crosswork Network Automation Functions in Real Time 128
Collect and Share Cisco Crosswork Network Automation Logs and Metrics 132
Control Cisco Crosswork Network Automation Applications and Services 133
Manage Backup and Restore 134
Disaster Restore 136
Configuration Database CLI Tool 137
Manage Users 137
Administrative Users Created During Installation 138
Add Users 138
Edit Users 139
Delete Users 139
Create User Roles 140
Edit User Roles 140
Clone User Roles 141
Delete User Roles 141
Manage TACACS+ Servers 142
Add a TACACS+ Server 142
Edit a TACACS+ Server 142

Delete a TACACS+ Server 143
Manage LDAP Servers 143
Add a LDAP Server 143
Edit a LDAP Server 145
Delete a LDAP Server 145
Define Network Topology Display Settings 145
Manage Certificates 145
Extend Self-Signed Certificate Expiration 147
Substitute a User-Provided Certificate 147
Smart Licensing Registration 148
Overview 148
Configure Transport Settings 149
Register 150
Manual Actions 152
License Authorization Statuses 153
Security Hardening Overview 154
Core Security Concepts 154
HTTPS 154
SSL Certificates 154
1-Way SSL Authentication 155
Disable Insecure Ports and Services 156
Harden Your Storage 157

CHAPTER 7 **Manage Cisco Crosswork Data Gateway** 159
Overview of Cisco Crosswork Data Gateway 159
Manage Cisco Crosswork Data Gateway Instances 159
Add a Cisco Crosswork Data Gateway Instance 164
Update Cisco Crosswork Data Gateway Instance Enrollment Settings 164
View Enrollment Details 165
Change the Administration State of a Cisco Crosswork Data Gateway Instance 167
De-enroll a Cisco Crosswork Data Gateway Instance 168
Attach a Device to a Cisco Crosswork Data Gateway Instance 169
Detach a Device From a Cisco Crosswork Data Gateway Instance 171
View Cisco Crosswork Data Gateway Instance Health 172

Configure Cisco Crosswork Data Gateway Settings 175
Manage Data Destinations 176
Add a Data Destination 177
Update a Data Destination 180
View Data Destination Details 180
Delete a Data Destination 181
Manage Custom Software Packages 182
Add a Custom Software Package 184
Delete a Custom Software Package 185
Download Custom or System MIBs and Packages 186

CHAPTER 8
Configure Collection 189
Collection Service Overview 189
Collection Considerations 189
Prerequisites for Device Model Driven Telemetry 190
About Collection Jobs 192
Collection Jobs 193
CLI Collection Job 193
SNMP Collection Jobs 194
MDT Collection Job 200
Monitoring Collection Jobs 201

APPENDIX A
Configure Cisco Crosswork Data Gateway Base VM 207
About Cisco Crosswork Data Gateway Base VM 207
Base VM Contents 207
Log In and Log Out 208
Access Cisco Crosswork Data Gateway Through vCenter 208
Access Cisco Crosswork Data Gateway Via SSH 209
Use the Interactive Console 209
Basic Concepts 210
Cisco Crosswork Data Gateway Components 210
Controller Gateway 211
Image Manager 211
Vitals Monitor 211

Route Manager 212
Docker IPv6nat 212
Manage Users 212
Supported User Roles 212
Change Password 214
View Current System Settings 214
Change Current System Settings 215
Configure NTP 217
Configure DNS 218
Configure Control Proxy 218
Configure Static Routes 218
Add Static Routes 218
Delete Static Routes 220
Configure Syslog 221
Create New SSH Keys 222
Import Certificate 222
Monitor Cisco Crosswork Data Gateway Health 222
Vitals Monitor 223
View Cisco Crosswork Data Gateway Vitals 223
collector-vitals Service 226
Troubleshooting 228
Ping a Host 229
Traceroute to a Host 230
Check NTP Status 231
Check System Uptime 231
Run show-tech 232
Reboot Crosswork Data Gateway VM 233

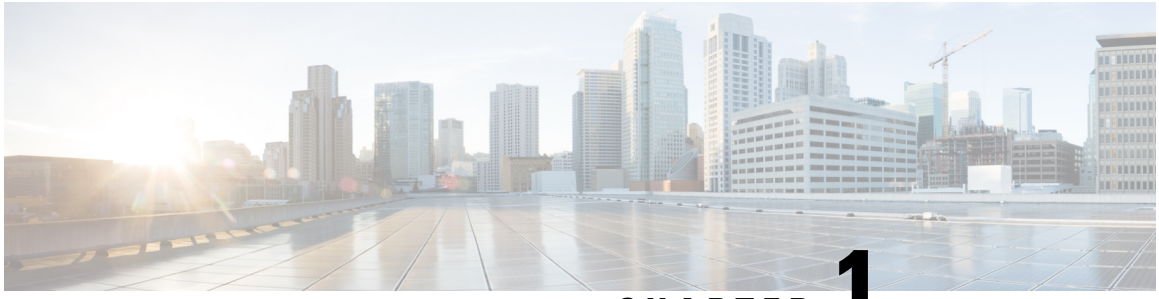

CHAPTER 1

Overview of Cisco Crosswork Optimization Engine

This section mainly describes what Cisco Crosswork Optimization Engine does and how to navigate the main user interface. To quickly get started, you should understand some basic concepts and look over the high-level workflows described in Get Started, on page 11.

The following topics are addressed in this section:

- Audience, on page 1
- Overview of Cisco Crosswork Optimization Engine, on page 2
- Cisco Crosswork Optimization Engine Architecture, on page 2
- Crosswork Optimization Engine APIs, on page 3
- Segment Routing Path Computation Element (SR-PCE), on page 3
- Log In and Log Out, on page 3
- Crosswork Optimization Engine Home Page, on page 4
- Set, Sort and Filter Table Data, on page 7

Audience

This guide is for experienced network administrators who want to use Cisco Crosswork Optimization Engine in their network. This guide assumes that you are experienced and familiar with using the following technologies:

- Networking technologies and protocols (BGP-LS, IGP (OSPF and IS-IS), PCEP, model-driven telemetry, and so on)
- Cisco IOS XR Traffic Controller (XTC) or Segment Routing Path Computation Element (SR-PCE) functionality
- RSVP-TE tunnel provisioning
- Segment routing (SR) policy provisioning

Overview of Cisco Crosswork Optimization Engine

Cisco Crosswork Optimization Engine is part of the Cisco Crosswork Network Automation suite of products. Cisco Crosswork Optimization Engine provides real-time network optimization allowing operators to effectively maximize network utilization as well as increase service velocity.

Crosswork Optimization Engine provides the following:

- A topology map that gives valuable real-time visualization of devices, links, link utilization, and SR policy or RSVP-TE tunnel provisioning in the network.

  To view supported TE tunnel features and limitations, see SR Policy and RSVP-TE Tunnel Support, on page 83.

- A UI that allows for easy manageability of SR policies and RSVP-TE tunnels. Crosswork Optimization Engine enables the network operator to perform the following tasks:
  - Provision SR policies and RSVP-TE tunnels and modify or remove them using an intuitive workflow
  - Continuously track SR policy dynamic path computations to maintain SLA objectives (with correct licensing)
  - Preview an SR policy or RSVP-TE tunnel before deploying it to the network
- APIs to extend Crosswork Optimization Engine functions. See the Cisco Crosswork Network Automation API Documentation on Cisco DevNet.
- Crosswork Optimization Engine function packs (with correct licensing) that provide closed-loop optimization to define the optimization intent, implement the intent, and continuously monitor, track, and react to maintain the original intent. See the *Cisco Crosswork Optimization Engine Function Packs* document.

Note To get a quick overview on how to start using Crosswork Optimization Engine, see High-Level Workflows, on page 15.

Cisco Crosswork Optimization Engine Architecture

In order to provide for better scalability and improved performance the data collection functionality has been separated out into its own VM and software package called Cisco Crosswork Data Gateway. The license to use Cisco Crosswork Data Gateway is included with the Crosswork Optimization Engine license. Cisco Crosswork Data Gateway gathers all the information from the monitored devices and forwards it to Crosswork Optimization Engine for analysis and processing. Crosswork Optimization Engine can then be used by the operator to manage the network or respond to changes in the network.

Apart from Crosswork Optimization Engine, Cisco Crosswork Data Gateway is required for external data collection, such as interface statistics via SNMP and model-driven telemetry sensor paths. Crosswork Optimization Engine can use Cisco Network Services Orchestrator (Cisco NSO) as a provider to manage the devices for any required model-driven telemetry (MDT) sensor paths for data collection. Cisco NSO then supplies the device management and configuration-maintenance services.

If you do not plan to use to use Cisco NSO, you must apply the telemetry configuration on your devices. See Prerequisites for Device Model Driven Telemetry, on page 190.

This guide explains how to use both Crosswork Optimization Engine and Cisco Crosswork Data Gateway.

For more information on configuring and managing Cisco Crosswork Data Gateway, see Manage Cisco Crosswork Data Gateway, on page 159.

Note Crosswork Optimization Engine is designed and tested to be used with the Cisco Crosswork Data Gateway 1.1 release.

Crosswork Optimization Engine APIs

Advanced users can extend Crosswork Optimization Engine functions by using product application programming interfaces (APIs).

For more information about the product APIs, see the Cisco Crosswork Network API Documentation on Cisco DevNet.

Disclaimer:

Cisco may be providing you with API software currently at no charge. However, nothing restricts Cisco's right, now or in the future, to monetize the API software. At which point you may be required to pay a license fee in order to use the API software.

Segment Routing Path Computation Element (SR-PCE)

Crosswork Optimization Engine uses the combination of telemetry and Cisco Segment Routing Path Computation Element (SR-PCE) to analyze and compute optimal TE tunnels.

Cisco SR-PCE (formerly Cisco XR Traffic Controller (XTC)) runs on the Cisco IOS XR operating system. SR-PCE provides stateful PCE functionality that helps control and move TE tunnels to optimize the network. PCE describes a set of procedures by which a Path Computation Client (PCC) can report and delegate control of head-end tunnels sourced from the PCC to a PCE peer. The PCC and PCE establish a Path Computation Element Communication Protocol (PCEP) connection that SR-PCE uses to push updates to the network.

Crosswork Optimization Engine discovers all devices that are part of the IGP domain including those that do not establish PCEP peering with SR-PCE. However, PCEP peering is required to deploy TE tunnels to the device.

Log In and Log Out

The Cisco Crosswork Optimization Engine user interface is browser based. See the *Cisco Crosswork Optimization Engine Installation Guide* for supported browser versions.

Step 1 Open a web browser and enter:

```
https://<CrossworkVMManagementIPAddress>:30603/
```

When you access Cisco Crosswork Optimization Engine from your browser for the first time, some browsers display a warning that the site is untrusted. When this happens, follow the prompts to add a security exception and download the self-signed certificate from the server. After you do this, the browser accepts the Cisco Crosswork Optimization Engine server as a trusted site in all subsequent logins.

Step 2 The Cisco Crosswork Optimization Engine browser-based user interface displays the login window. Enter your username and password.

Note The default Cisco Crosswork Optimization Engine administrator user name and password is **admin**. This account is created automatically at installation (see Administrative Users Created During Installation, on page 138). The initial password for this account must be changed during installation verification. Cisco strongly recommends that you keep the default administrator credential secure, and never use it for routine logins. Instead, create new user accounts with appropriate privileges and their own credentials (as explained in Add Users, on page 138) and use only those accounts for all subsequent user logins.

Step 3 Click **Log In**.

Step 4 To log out, click in the top right of the Cisco Crosswork Optimization Engine main window and choose **Log out**.

Crosswork Optimization Engine Home Page

Figure 1: Crosswork Optimization Engine Home Page

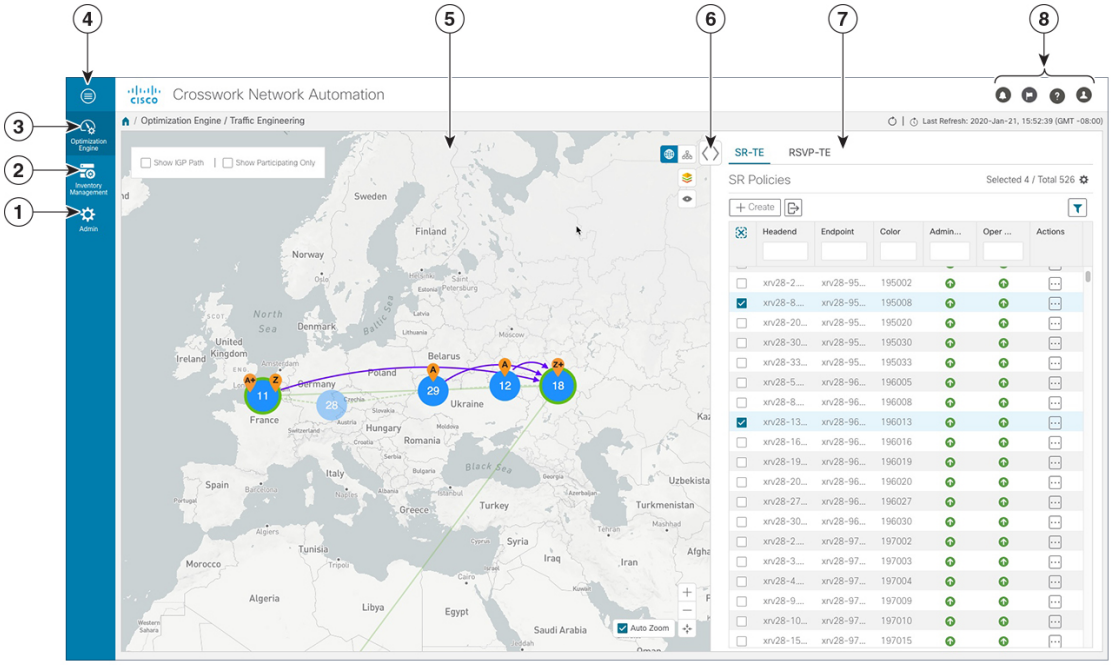

Callout No.	Description
1	**More**: Toggles the main menu to compact mode or expanded mode. In compact mode, you must hover over the main menu items to view and select available options. In expanded mode, you must click on the main menu item to display the available options. In this mode, when a main menu item is expanded, it will remain so until you collapse the menu item.
2	**Network Topology Map**: Displays a geographical or logical map view of the devices, links, and SR policies in your network. It also shows the general condition of devices and links. See Visualize the Network, on page 65. In conjunction with the SR Policies Table and RSVP-TE Tunnels Table, it quickly highlights selected TE tunnels and associated tunnel information such as metrics, adjacency segment IDs, segment hops, source and destination nodes. See Visualize SR Policies and RSVP-TE Tunnels, on page 83.
3	**Expand/Collapse/Hide Side Panel**: Expand or collapse the contents of the side panel. Close the side panel to get a larger view of the topology map.
4	The content of this panel changes depending if the SR-TE tab (SR Policies Table, on page 94) or RSVP-TE tab (RSVP-TE Tunnels Table, on page 96) is selected. Depending on what is selected on the topology map, or whether you are in the process of viewing and managing TE tunnels, you can do the following: • Create and Manage SR Policies, on page 107 • Create and Manage RSVP-TE Tunnels, on page 117 • Get More Information About an SR Policy, on page 114 • Get More Information About an RSVP-TE Tunnel, on page 121 • Get More Information About Devices on the Map, on page 74 • Get More Information About Links, on page 78
5	Settings icons: The **Alerts** icon notifies you of any current error conditions related to the system operations which require attention, and provides a link to detailed information about those conditions. The **Events** icon notifies you of new events related to system operation, and also provides access to the history of all system events. The **About** icon displays the current version of Crosswork Optimization Engine. The **User Account** icon lets you view your username, change your password, and log out.

Callout No.	Description
6	**Optimization Engine Menu**: You can access the following TE tunnel related options: • **Traffic Engineering**—Returns you to the main window as shown above. • **Affinity Mapping**—Lets you map an affinity to a bit position. See Configure Affinity Mapping, on page 103. • **Function Packs**—Lets you enable and configure function packs. See the *Cisco Crosswork Optimization Engine Function Packs* document.
7	**Inventory Management Menu**: You can access the following inventory related options: • **Devices**—Lets you add, delete, update and view information about the devices in your network. See Manage Network Devices, on page 48. • **Providers**—Lets you add, delete, update and manage providers. See Manage Providers, on page 32. • **Credentials**—Lets you add, delete, update and manage credential profiles that control access to devices and providers. See Manage Credential Profiles, on page 25. • **Tags**—Lets you add, delete, update and manage the tags you use to sort and group devices. See Manage Tags, on page 59. • **Job History**—Lets you review device related jobs. See View Device Job History, on page 58.

Callout No.	Description

8	**Admin Menu**: You can access the following administrative related options:
	• **Crosswork Manager**—Lets you do the following tasks:
	  • Collect logs and metrics. See Collect and Share Cisco Crosswork Network Automation Logs and Metrics, on page 132.
	  • Monitor the general state of containers. See Monitor Cisco Crosswork Network Automation Functions in Real Time, on page 128.
	  • Control (stop, start, or restart) services. See Control Cisco Crosswork Network Automation Applications and Services, on page 133.
	• **Backup Restore**—Lets you restore or create a backup for the Crosswork Optimization Engine VM.
	• **Users**—Lets you add, update and view users and roles. See Manage Users, on page 137.
	• **AAA**—Lets you add, update and view TACACS+ and LDAP to authenticate users. See Manage TACACS+ Servers, on page 142.
	• **Visualization Settings**—Lets you update topology map settings. See Configure Geographical Map Settings, on page 69 and Define Color Thresholds for Link Bandwidth Utilization, on page 81.
	• **Certificate Management**—Lets you view and manage certificates. See Manage Certificates, on page 145.
	• **Collection Jobs**—Lets you view collection job details. See Monitoring Collection Jobs, on page 201.
	• **Data Gateway Management**—Lets you manage Cisco Crosswork Data Gateway instances and view their metrics. See Manage Cisco Crosswork Data Gateway, on page 159.
	• **Data Gateway Global Settings**—Lets you set up data collection for external data destinations and add custom software packages (custom MIB packages, CLI device package, and SNMP device packages). See Configure Cisco Crosswork Data Gateway Settings, on page 175.
	• **Smart Licensing Registration**—Lets you register, view, and manage your Smart Licenses. See Smart Licensing Registration, on page 148.

Set, Sort and Filter Table Data

Many Cisco Crosswork Optimization Engine windows show database records in tables.

Any window with a table will also provide column selection, sorting, and filter functions that let you control the database records shown in the tables and help you locate particular records quickly.

Click ⚙ to display a list of all the fields in the database for the kind of data record displayed in the table. You can choose which fields you want to display as table columns by checking or unchecking the box next to any field in the list. Your choices are enabled immediately and are permanent.

You can also sort all the records displayed in the table according to the data in any one column by clicking that column's title:

- To sort the records in ascending order, click the column title once.
- To sort the records in descending order, click the column title again.

Sorting takes place immediately. You can only have one active sort at a time. The example **Links** window, below, shows an active sort on the **Link Type** field.

You can also filter the table to show only the records you want, using a quick filter . Many tables have all these features enabled by default. If you cannot see the quick filter features displayed on a window with a table, click .

The quick filter displays only the records that match the value you enter above the column in the **quick filter** field (see item 2, below). Filtering takes place immediately, as you type.

The advanced filter (only available in some tables) narrows the content in the table by applying a filter that includes both a value and a logical operator, such as Equals, Starts with, Contains, and so on. Click in the column header to access the advanced filter (see items 4 and 5, below).

In addition to these quick and advanced filters, you can also use tags to filter the devices shown in the **Devices** window (see Filter Devices by Tags, on page 56).

Figure 2: Links Window With Active Sort and Filters

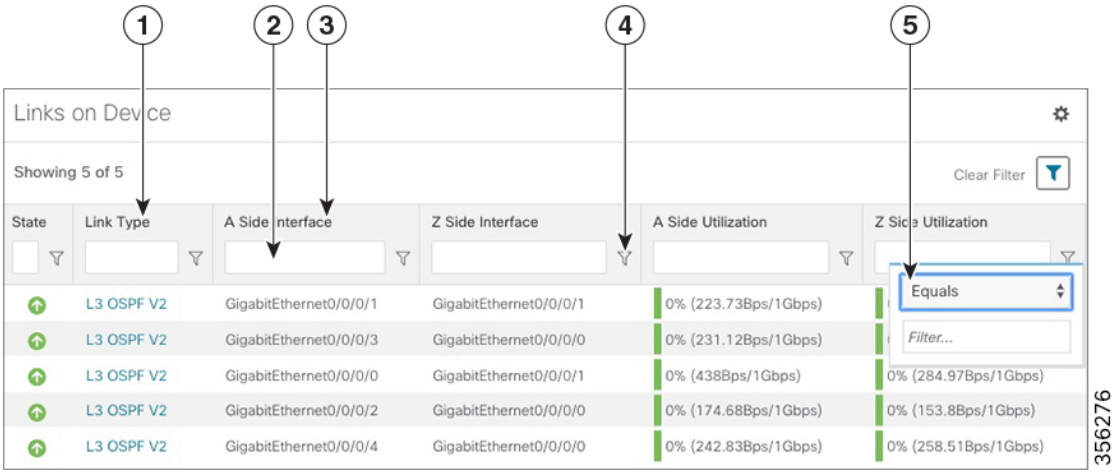

Item	Description
1	**Sort active icon**: This arrow icon indicates that the user has sorted the links by clicking on the column header. The arrow's direction shows that the table is sorted by **Link Type**, in ascending order.
2	**Quick filter field**: Type a text or numeric value in this field to show only the links that match the value you enter. The field shows the values you entered for both quick and advanced filters.
3	**Filter active icon**: This icon shows that a quick or advanced filter is currently applied to the data in this column.

Item	Description
4	**Advanced filter icon**: Click , shown in each column header, to specify an advanced filter on that column, using logical operators as well as alphanumerical values. **Note** Advanced filtering is not available on all tables.
5	**Filter criteria fields**: These fields appear in a popup next to the column after you click the icon. Set the filter criteria by selecting the logical operator from the drop down list in the first field, and then entering the filter value in the second field. Your criteria will be applied immediately. You will then be prompted to enter more operators and values, and to decide if you want to concatenate them using logical AND or OR. The quick filter field shows the values you entered (but not the operators). Logical operators include Equals, Not equal, Starts with, Ends with, Contains, and Not contains. **Note** Some columns will not have all of the logical operators available.

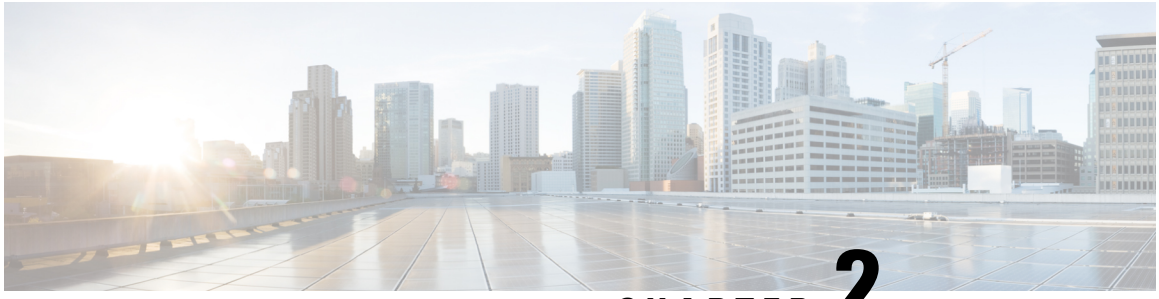

CHAPTER 2

Get Started

Resource Reservation Protocol (RSVP) is a protocol that you are most likely familiar with. It is a signaling protocol that enables systems to request resource reservations from the network. RSVP processes protocol messages from other systems, processes resource requests from local clients, and generates protocol messages. As a result, resources are reserved for data flows on behalf of local and remote clients. RSVP creates, maintains, and deletes these resource reservations. This chapter focuses on a high-level overview of Segment Routing (SR) which is gaining popularity in the networking routing area, as well as quick start workflows to help you get started with Crosswork Optimization Engine.

- Segment Routing, on page 11
- Inventory Management Concepts, on page 13
- Before You Begin, on page 14
- High-Level Workflows, on page 15

Segment Routing

Segment routing is a method of forwarding packets on the network based on the source routing paradigm. The source chooses a path and encodes it in the packet header as an ordered list of segments. Segments are an identifier for any type of instruction. For example, topology segments identify the next hop toward a destination. Each segment is identified by the segment ID (SID) consisting of a flat unsigned 32-bit integer.

Segments

Interior gateway protocol (IGP) distributes two types of segments: prefix segments and adjacency segments. Each router (node) and each link (adjacency) has an associated segment identifier (SID).

- A prefix SID is associated with an IP prefix. The prefix SID is manually configured from the segment routing global block (SRGB) range of labels, and is distributed by IS-IS or OSPF. The prefix segment steers the traffic along the shortest path to its destination. A node SID is a special type of prefix SID that identifies a specific node. It is configured under the loopback interface with the loopback address of the node as the prefix.

  A prefix segment is a global segment, so a prefix SID is globally unique within the segment routing domain.

- An adjacency segment is identified by a label called an adjacency SID, which represents a specific adjacency, such as egress interface, to a neighboring router. The adjacency SID is distributed by IS-IS or OSPF. The adjacency segment steers the traffic to a specific adjacency.

An adjacency segment is a local segment, so the adjacency SID is locally unique relative to a specific router.

By combining prefix (node) and adjacency segment IDs in an ordered list, any path within a network can be constructed. At each hop, the top segment is used to identify the next hop. Segments are stacked in order at the top of the packet header. When the top segment contains the identity of another node, the receiving node uses equal cost multipaths (ECMP) to move the packet to the next hop. When the identity is that of the receiving node, the node pops the top segment and performs the task required by the next segment.

Segment Routing for Traffic Engineering

Segment routing for traffic engineering takes place through a tunnel between a source and destination pair. Segment routing for traffic engineering uses the concept of source routing, where the source calculates the path and encodes it in the packet header as a segment. Each segment is an end-to-end path from the source to the destination, and instructs the routers in the provider core network to follow the specified path instead of the shortest path calculated by the IGP. The destination is unaware of the presence of the tunnel.

Segment Routing Policies

Segment routing for traffic engineering uses a “policy” to steer traffic through the network. An SR policy path is expressed as a list of segments that specifies the path, called a segment ID (SID) list. Each segment is an end-to-end path from the source to the destination, and instructs the routers in the network to follow the specified path instead of the shortest path calculated by the IGP. If a packet is steered into an SR policy, the SID list is pushed on the packet by the head-end. The rest of the network executes the instructions embedded in the SID list.

Note Cisco Crosswork Optimization Engine discovers existing SR policies when devices are imported, but cannot manage them. SR policies can be managed only if they were provisioned using Cisco Crosswork Optimization Engine (see Create and Manage SR Policies, on page 107).

There are two types of SR policies: dynamic and explicit.

Dynamic SR Policy

A dynamic path is based on an optimization objective and a set of constraints. The head-end computes a solution, resulting in a SID list or a set of SID lists. When the topology changes, a new path is computed. If the head-end does not have enough information about the topology, the head-end might delegate the computation to a path computation engine (PCE).

Explicit SR Policy

When you configure an explicit policy, you specify an explicit path which consists of a list of prefix or adjacency SIDs, each representing a node or link along on the path.

Disjointness

Cisco Crosswork Optimization Engine uses the disjoint policy to compute two lists of segments that steer traffic from two source nodes to two destination nodes along disjoint paths. The disjoint paths can originate from the same head-end or different head-ends. Disjoint level refers to the type of resources that should not be shared by the two computed paths. The following disjoint path computations are supported:

- **Link** – Specifies that links are not shared on the computed paths.

- **Node** – Specifies that nodes are not shared on the computed paths.
- **SRLG** – Specifies that links with the same Share Risk Link Group (SRLG) value are not shared on the computed paths.
- **SRLG-node** – Specifies that SRLG and nodes are not shared on the computed paths.

When the first request is received with a given disjoint-group ID, a list of segments is computed, encoding the shortest path from the first source to the first destination. When the second request is received with the same disjoint-group ID, information received in both requests is used to compute two disjoint paths: one path from the first source to the first destination, and another path from the second source to the second destination. Both paths are computed at the same time. The shortest lists of segments is calculated to steer traffic on the computed paths.

Note

- Disjointness is supported for two policies with the same disjoint ID.
- Configuring affinity and disjointness at the same time is not supported.

Inventory Management Concepts

Crosswork Optimization Engine makes extensive use of three basic inventory management concepts. It is helpful to be familiar with them before you get started.

- **Tags**: Tags will be familiar from other Web applications. They are simple text strings you can attach to objects to help group them. Crosswork Optimization Engine comes with a short list of ready-made tags used to group network devices. You can create your own tags and use them to identify, find, and group devices for a variety of purposes. For example, in addition to type and geolocation, you may want to identify and group them by their location in your network topology (Spine vs. Leaf), or the function they serve on your network (Provider vs. ProviderEdge). You will want to develop your own tags for your purposes, and rework them as needed to meet changing needs.
- **Providers**: Crosswork Optimization Engine does not perform inventory collection, route segmentation or configuration changes directly. Instead, it relies on an SR-PCE provider to perform these functions. The provider family determines the type of service that provider supplies to Crosswork Optimization Engine, and the parameters unique to that service, which must be configured. This architecture permits Crosswork Optimization Engine to devote all of its resources to processing and interpreting network events and rolling out changes in response to these events.
- **Credential Profiles:** For Crosswork Optimization Engine to be able to access a device or to interact with a provider, it must be able to present credentials. Rather than entering credentials each time they are needed, you can instead create credential profiles to securely store this information. The platform supports unique credentials for each type of access protocol, and allows you to bundle multiple protocols and their corresponding credentials in a single profile. Devices that use the same credentials can share a credential profile. For example, if all of your routers in a particular building share a single SSH user ID and password, you can create a single credential profile to allow Crosswork Optimization Engine to access and manage them.

Before You Begin

Before you begin using Cisco Crosswork Optimization Engine, Cisco recommends that you complete the following planning and information-gathering steps, in any order you wish:

- **User Accounts** : Cisco recommends as a best practice that you create separate accounts for all of your users, so that there is an audit record of user activity on the system. Prepare a list of the people who will use Cisco Crosswork Optimization Engine. Decide on their user names and preliminary passwords, and create user profiles for them (see Manage Users, on page 137).
- **User Roles**: Cisco recommends that you use role-based access control to confine users to just the software functions needed to perform their job duties. By default, every new user you create has full administrative privileges. Unless you want to extend the same privileges to every user, you will need to plan a system of user roles, create them, and assign them to the user profiles you create (see Create User Roles, on page 140).
- **Credentials**: Gather access credentials and supported protocols that you will use to monitor and manage your devices. For providers, this always includes user IDs, passwords, and connection protocols. For devices, it includes user IDs, passwords, and additional data such as the SNMP v2 read and write community strings, and SNMPv3 auth and privilege types. You will use these to create credential profiles (see Inventory Management Concepts, on page 13 and Manage Credential Profiles, on page 25).
- **Tags**: Plan a preliminary list of custom tags to create when setting up the system, so that you can use them to group your devices when you first onboard them. As explained in Inventory Management Concepts, on page 13, you will want to consider grouping devices by functionality. You need not have a complete list of tags at first, as you can always add more later, but please note that all the tags you do plan to use must be in place before you need them; you cannot create them "on the fly" (see Manage Tags, on page 59 and Create Tags, on page 60).
- **Providers**: As explained in Inventory Management Concepts, on page 13, providers do the basic work of direct interaction with network devices, so that Cisco Crosswork Optimization Engine can automate monitoring and responses to network events. At a minimum, Cisco Crosswork Optimization Engine must have an SR-PCE provider defined in order to discover devices and to distribute policy configuration to devices. You should determine the auto-onboarding mode and device profile you will use (if you auto-onboard devices). See Add Cisco SR-PCE Providers, on page 34.
- **Devices**: Decide how you are going to onboard your devices: manually, via the user interface, or automatically via synchronization or CSV import. This determines the amount of additional information you will need to onboard your devices, which is covered in About Adding Devices, on page 19.
- **External Data Destination(s)**: Decide which external data destination (Kafka or gRPC) you are going to use and ensure it is set up to receive input from Cisco Crosswork Data Gateway.

Note that you can capture the devices, credential profiles, tags, and providers lists in spreadsheet form, convert the spreadsheet to CSV format, and then upload them in bulk to Cisco Crosswork Optimization Engine. You do this using the Import feature (accessed using the Import icon,).

You can access CSV templates for each of these lists by clicking the Import icon in the corresponding places in the user interface. Select the **Download template** link when prompted to choose an export destination path and file name.

High-Level Workflows

These workflows describe the steps to quickly get started with Cisco Crosswork Optimization Engine. The main difference between the two workflows are the steps on how devices are added (see About Adding Devices, on page 19). You will find that the general order of steps (adding a Cisco Crosswork Data Gateway instance, creating device tags, etc.) can be done out of the documented order as long as you keep the following in mind:

- A Cisco Crosswork Data Gateway instance must be created.
- In order for data collection to occur, devices must be attached to a Cisco Crosswork Data Gateway instance.
- An SR-PCE provider must be configured before devices are added.

Note If you selected to use Cisco NSO for device management during Cisco Crosswork Optimization Engine installation, you must add NSO as a provider (see *Cisco Crosswork Optimization Engine Installation Guide*).

Workflow: Auto-Onboard Devices

The following workflow describes the main steps to get started with Cisco Crosswork Optimization Engine by configuring a Cisco SR-PCE provider to automatically onboard devices.

Table 1: Workflow: Automatic Onboarding of SR-PCE Devices

Step	For more information, see...
1. Ensure that your devices are configured properly for communication and telemetry.	Refer to the guidelines and sample configurations in: • Prerequisites for Onboarding Devices, on page 21 • Sample Configuration for Devices in Cisco NSO, on page 22 **Note** Only if NSO is being used for device management. • Prerequisites for Device Model Driven Telemetry, on page 190
2. Create a device credential profile.	Create Credential Profiles, on page 26

Step	For more information, see...
3. Configure SR-PCE as a provider. **Note** The **auto-onboard** provider property value must be set to **managed** or **unmanaged** to enable automatic onboarding of devices. For more information see About Adding Devices, on page 19 and Auto-Onboard Property Descriptions, on page 36.	Add Cisco SR-PCE Providers, on page 34
4. Validate communications with provider.	Get Provider Details, on page 45
5. (Required if using NSO for device management) Configure NSO credential profile and provider.	• Create Credential Profiles, on page 26 • Add Cisco NSO Providers, on page 42
6. (Required if using NSO for device management, otherwise optional) To update device attributes (such as mapping a device to NSO, adding connectivity IP and geographical coordinates, and so on) export the CSV device list, make and save modifications, and import it back to the device inventory. **Note** If you wish to use the geographical topology map, you must add geographical location details.	• Export Devices, on page 57 • Import Devices, on page 49
7. (Optional) Create tags for use in grouping new devices.	Manage Tags, on page 59
8. Add a Cisco Crosswork Data Gateway instance. At least one instance should have been added and enrolled during installation.	See the *Cisco Crosswork Optimization Engine Installation Guide*.
9. Confirm that the Cisco Crosswork Data Gateway instance is running. The operational state of the Cisco Crosswork Data Gateway instance must be **Up** before continuing to the next step.	Manage Cisco Crosswork Data Gateway Instances, on page 159
10. Attach devices to Cisco Crosswork Data Gateway so that data can be collected.	Attach a Device to a Cisco Crosswork Data Gateway Instance, on page 169
11. View device list (**Inventory Management** > **Devices**) to check that devices have been added properly. If devices are unreachable, select and edit the device with connectivity details.	Manage Network Devices, on page 48 Click to investigate any device whose **Reachability State** is marked as (unreachable), (degraded), or (unknown).
12. Confirm visualization of IGP topology (logical view).	Network Topology Map, on page 65
13. Visualize discovered SR policies and RSVP-TE tunnels and create new tunnels.	Visualize SR Policies and RSVP-TE Tunnels, on page 83

Workflow: Manually Import Devices

The following workflow describes the main steps to get started with Cisco Crosswork Optimization Engine by importing a CSV file to add devices.

Table 2: Workflow: Importing a CSV file to Onboard Devices

Step	For more information, see...
1. Ensure that your devices are configured properly for communication and telemetry.	Refer to the guidelines and sample configurations in: • Prerequisites for Onboarding Devices, on page 21 • Sample Configuration for Devices in Cisco NSO, on page 22 **Note** Only if NSO is being used for device management. • Prerequisites for Device Model Driven Telemetry, on page 190
2. Create a device credential profile.	Create Credential Profiles, on page 26
3. Configure the SR-PCE provider. **Note** Set **auto-onboard** property value to **off** for manual device onboarding. For more information see Auto-Onboard Property Descriptions, on page 36.	Add Cisco SR-PCE Providers, on page 34
4. (Required if using NSO for device management) Configure NSO credential profile and provider.	• Create Credential Profiles, on page 26 • Add Cisco NSO Providers, on page 42
5. (Optional) Create tags for use in grouping new devices.	Manage Tags, on page 59
6. Create a CSV file and import devices.	Import Devices, on page 49
7. (Optional) Modify device details.	Edit Devices, on page 56
8. Add a Cisco Crosswork Data Gateway instance. At least one instance should have been added and enrolled during installation.	See the *Cisco Crosswork Optimization Engine Installation Guide*.
9. Confirm that the Cisco Crosswork Data Gateway instance is running. The operational state of the Cisco Crosswork Data Gateway instance must be **Up** before continuing to the next step.	Manage Cisco Crosswork Data Gateway Instances, on page 159
10. Attach devices to Cisco Crosswork Data Gateway so that data can be collected.	Attach a Device to a Cisco Crosswork Data Gateway Instance, on page 169

Step	For more information, see...
11. View device list (**Inventory Management** > **Devices**) to check that devices have been added properly. If devices are unreachable, select and edit the device with connectivity details.	Manage Network Devices, on page 48 Click to investigate any device whose **Reachability State** is marked as (unreachable), (degraded), or (unknown).
12. Confirm visualization of IGP topology (logical view).	Network Topology Map, on page 65
13. Visualize discovered SR policies and RSVP-TE tunnels and create new tunnels.	Visualize SR Policies and RSVP-TE Tunnels, on page 83

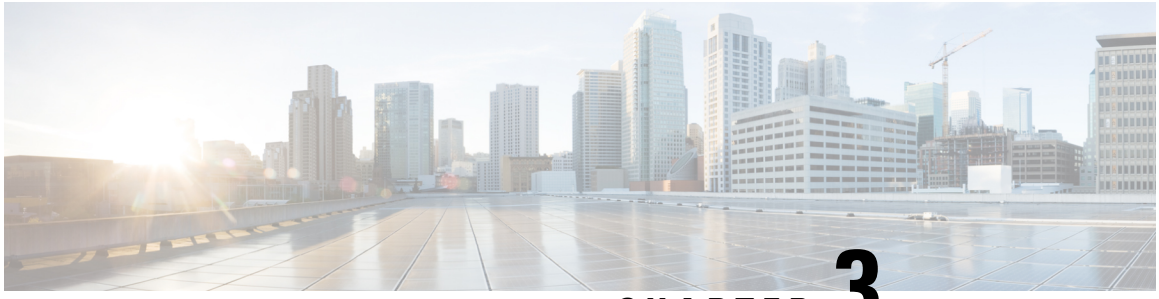

CHAPTER 3

Manage Inventory

This section contains the following topics:

- Device Management Overview, on page 19
- About Adding Devices, on page 19
- Prerequisites for Onboarding Devices, on page 21
- Sample Configuration for Devices in Cisco NSO, on page 22
- Reachability and Operational State, on page 23
- Manage Credential Profiles, on page 25
- Manage Providers, on page 32
- Manage Network Devices, on page 48
- Manage Tags, on page 59

Device Management Overview

The Device Management application lets you create, edit, and delete:

- The **credential profiles** that control Crosswork Optimization Engine's access to devices and providers. See Manage Credential Profiles, on page 25.
- The **devices** you manage using Crosswork Optimization Engine. See Manage Network Devices, on page 48.

You can also use Device Management to review the **jobs** executed on your devices. See View Device Job History, on page 58.

About Adding Devices

There are two ways to add devices to Cisco Crosswork Optimization Engine:

1. Automatically onboard devices and populate the inventory. High-level steps are documented in Workflow: Auto-Onboard Devices, on page 15.
2. Manually onboard devices using a CSV file or the UI. High-level steps are documented in Workflow: Manually Import Devices, on page 17

Auto-Onboard Devices

Auto-onboarding simplifies and expedites the device onboarding process. It automatically discovers and imports preformatted device data from a Cisco SR-PCE provider and enables you to quickly view the IGP topology (including devices, links and IP addresses) in the Cisco Crosswork Optimization Engine topology map.

To configure auto-onboarding, you add an SR-PCE provider with one of the following auto-onboard options: **`managed`** or **`unmanaged`**.

The auto-onboard **`managed`** option requires a single default credential profile (having SNMP access, at minimum) that will work for all devices.

The devices are auto-onboarded with the following attributes:

- **ISIS-System ID**, **OSPF Router ID**, and **TE router ID** will be filled in the device's routing information.
- The **Connectivity IP** is assigned the same value as the **TE router ID**.
- The default credential profile is set as the **Credential Profile** for each device.

Note If a common credential profile cannot be used for all devices, or a different **Connectivity IP** is required, use the auto-onboard **`unmanaged`** option or Cisco Crosswork Optimization Engine will keep trying to connect to the devices and fail.

The auto-onboard **`unmanaged`** option should be used if you prefer devices not to be assigned a **Credential Profile** or **Connectivity IP**. SNMP or any other device collection is not performed. However, IGP topology is still seen on the topology map (logical view), but the information available is restricted to the information SR-PCE provides. Therefore, interface names are not shown, and in the case of OSPF, device Hostnames are also not shown. IP addresses are shown and can be used to identify devices and interfaces.

Auto-Onboard Notes and Limitations:

Consider the following information when choosing between **`unmanaged`** and **`managed`** options:

- The TE router ID is used as the Connectivity IP of the device. This is the IP address Cisco Crosswork Optimization Engine will use to perform SNMP or CLI collection from the device. If the devices need to be reached over a separate management network, the Connectivity IP of all devices will need to be updated using the CSV **Update Existing** option (see Import Devices, on page 49). In this case, use the **`unmanaged`** option for auto-onboarding to prevent repeated unsuccessful collection attempts from the devices.
- The **`managed`** option works only if a single **Credential Profile** will work for accessing all the devices.
- With the **`unmanaged`** option, since SNMP collection from the devices cannot be performed, interface names and possibly hostnames will not be available until the devices in inventory are updated with the correct **Connectivity IP** and **Credential Profile** and their state is updated to Managed.
- Several device attributes cannot be discovered and need to be manually supplied. After the inventory is populated, you can download the device inventory CSV file, edit the file to add additional information (such as geographical location), and import it back into Cisco Crosswork Optimization Engine using the CSV **Update Existing** option. See Import Devices, on page 49 and Export Devices, on page 57.

Manually Add Devices

You can manually onboard devices from a CSV file or add them using the UI. After adding credential profiles, configure providers and tags to group new devices (optional) you do one of the following:

- Download the CSV template file from **Inventory Management** > **Devices** > and populate it with all the devices you will need (see Import Devices, on page 49). This method can be time consuming, as you must create and enter all of the data yourself beforehand (including not only devices, but also the providers, credential profiles and tags), and then ensure all of these items are properly associated with the devices.

  To quickly get up and running with Cisco Crosswork Optimization Engine by importing devices, follow the high-level steps documented in Workflow: Manually Import Devices, on page 17.

- Add devices using the UI (see Add Devices Through the UI, on page 51. It is the most time-consuming since all data is validated during entry.

Prerequisites for Onboarding Devices

Before adding devices, you must ensure that the devices themselves are configured to collect and transmit telemetry data properly and communicate successfully with Cisco Crosswork Optimization Engine. The following sections provide sample configurations for a variety of communications options. Use them as a guide to configuring the devices you plan to manage using Cisco Crosswork Optimization Engine.

Note Only users configured with privilege level 15 can use the NETCONF APIs. Privilege level 15 can be used to configure the "enable" password option in XE devices. In such cases, NETCONF should not be included as one of the protocols to verify reachability and operational state for the onboarded devices.

Note Only SNMPv2 and SNMPv3 (NoAuth/NoPriv) traps are supported.

Pre-Onboarding SNMP v2 Device Configuration

The following commands provide a sample pre-onboarding device configuration that sets the correct SNMPv2 and NETCONF configuration, and SSH and Telnet rate limits. The NETCONF setting is only needed if the device is MDT-capable (XR 6.5.3/6.6.3 or higher).

```
logging console debugging
logging monitor debugging
telnet vrf default ipv4 server max-servers 100
telnet vrf default ipv6 server max-servers 100
crypto key generate rsa
line default
 exec-timeout 0 0
 width 107
 length 37
 absolute-timeout 0
!
snmp-server community public RO
snmp-server community robot-demo2 RO
snmp-server ifindex persist
ntp
```

```
 server <NTPServerIPAddress>
!
service cli history size 5000
service cli interactive disable
ssh server v2
ssh server vrf default
ssh server netconf vrf default
ssh server logging
ssh server rate-limit 100
ssh server session-limit 100
grpc
 port 57400
!
netconf agent tty
!
netconf-yang agent
 ssh
!
```

Pre-Onboarding SNMPv3 Device Configuration

If you want to enable SNMPv3 data collection, repeat the SNMPv2 configuration commands in the previous section, and add the following commands:

```
snmp-server group grpauthpriv v3 priv notify v1default
snmp-server user <user-ID> grpauthpriv v3 auth md5 <password> priv aes 128 <password>
```

Sample Configuration for Devices in Cisco NSO

If you plan to use Cisco NSO as a provider to configure devices managed by Cisco Crosswork Optimization Engine, be sure that the Cisco NSO device configurations observe the following guidelines.

The following example shows a Cisco NSO setup that uses the hostname as the device ID. If you are using a CSV file to import devices, use **ROBOT_PROVDEVKEY_HOST_NAME** as the enum value for the provider_node_key field. The example hostname **RouterFremont** used here must match the hostname for the device in the CSV file.

```
configure
set devices device RouterFremont address 198.18.1.11 port 22
set devices device RouterSFO address 198.18.1.12 port 830
```

The authgroup username and password in the CSV file must match the username and password in the credential profile associated with the Cisco NSO provider. For example:

```
set devices authgroups group cisco default-map remote-name cisco remote-password cisco
set devices device Router* device-type netconf ned-id cisco-iosxr-nc-6.6
set devices device Router* authgroup cisco
```

The device itself must be synchronized with Cisco NSO before you import that device. For example:

```
set devices device Router* state admin-state unlocked
request devices device Router* ssh fetch-host-keys
request devices device Router* sync-from
commit
```

Reachability and Operational State

Cisco Crosswork Optimization Engine computes the Reachability State of the providers it uses and devices it manages, as well as the Operational State of reachable managed devices. It indicates these states using the icons in the following table.

Table 3: Reachability and Operational State Icons

This Icon...	Indicates...
Reachability State icons show whether a device or a provider is reachable or not	
	Reachable: The device or provider can be reached by all configured protocols configured for it.
	Reachability Degraded: The device or provider can be reached by at least one protocol, but is not reachable by one or more of the other protocols configured for it.
	Unreachable: The device or provider cannot be reached by reachable by any protocol configured for it.
	Reachability Unknown: Cisco Crosswork Optimization Engine cannot determine if the device is reachable, degraded, or unreachable . This state can also occur if the device is not connected to Cisco Crosswork Data Gateway.
Operational State icons show whether a device is operational or not.	
	The device is operational and under management, and all individual protocols are "OK" (also known as "up").
	The device is not operational ("down"). The same icon is used when the device has been set "administratively down" by an operator.
	The device's operational or configuration state is unknown.
	The device's operational or configuration state is degraded.

This Icon...	Indicates...
	The device's operational or configuration state is in an error condition. It is either not up, or unreachable, or both, due to errors encountered while attempting to reach it and compute its operational state. The number in the circle shown next to the icon indicates the number of recent errors. Click on the number to see a list of these errors. (Note that the icon badging for errors is not available in the Network Topology application.)
	The device's operational state is currently being checked
	The device is being deleted.
	The device is unmanaged.

The Reachability State of a device is computed as follows:

1. Reachability is always computed for each device as long as the device's configured state (as configured by users) is UP. It is not computed if the device is administratively DOWN or UNMANAGED.
2. Reachability state is always either REACHABLE, UNREACHABLE, or UNKNOWN.
   - The Reachability state is REACHABLE if there is at least one route to the device via at least one protocol AND the device is discoverable.
   - The Reachability state is UNREACHABLE if there are no routes to the device via one protocol OR the device does not respond.
   - The Reachability state is UNKNOWN if the device is UNMANAGED.

The Operational State of a device is computed as follows:

1. Operational state is always computed for each device as long as the device's configured state (as configured by users) is UP. It is not computed if the device is administratively DOWN or UNMANAGED.
2. Operational state is always OK or ERROR.
3. For a device to be Operational=OK, the device must be REACHABLE and discoverable. Any other Reachability state is ERROR.
4. For XR or XE devices only, Operational=OK also requires that Clock Drift difference between the Crosswork host and device clocks is <=the default Drift Value, currently 2 minutes.

Note Confirm that devices have Telnet/SSH enabled. If it is not enabled, the Clock Drift throws an error and the operational state will always show a clock synchronization error.

Manage Credential Profiles

Credential profiles are collections of credentials for SNMP, Telnet/SSH, HTTP, and other network protocols. You can have multiple protocols and credentials in a single credential profile.

Using credential profiles lets you automate device configuration changes and monitoring, and communicate with providers. When you add or import devices, or create providers, you specify the credential profile(s) those devices and providers use.

Note Credentials just validates authentication since the corresponding protocol configured on the devices does the work. Devices should be present in the **Devices** window and be reachable.

From the **Credential Profiles** window, you can create a new credential profile, update the settings configured for an existing profile, or delete a profile. To open this window, choose **Inventory Management** > **Credential Profiles** from the main menu.

Figure 3: Credentials Profile window

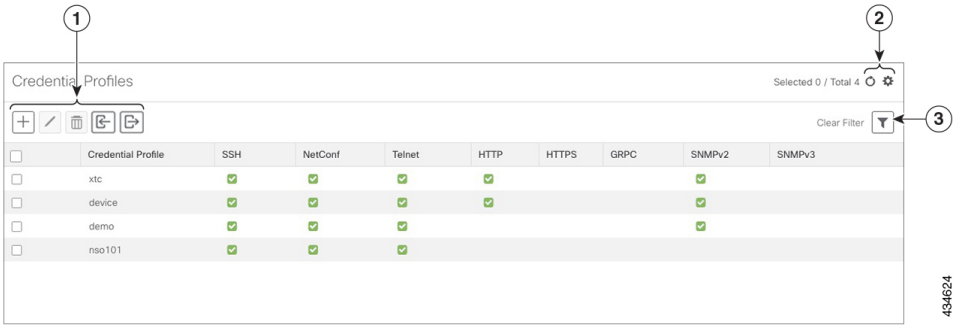

Item	Description
1	Click + to add a credential profile. See Create Credential Profiles, on page 26.
	Click ✎ to edit the settings for the selected credential profile. See Edit Credential Profiles, on page 30.
	Click 🗑 to delete the selected credential profile. See Delete Credential Profiles, on page 30.
	Click ⇦ to import new credential profiles from a CSV file. You can also download a CSV file template by clicking this icon. The template includes sample data that you can use as a guide for building your own CSV file. See Import Credential Profiles, on page 27.
	Click ⇨ to export credential profiles to a CSV file. See Export Credential Profiles, on page 30.

Item	Description
2	Click to refresh the **Credential Profiles** window.
	Click to choose the columns to make visible in the **Credential Profiles** window (see Set, Sort and Filter Table Data, on page 7).
3	Click to set filter criteria on one or more columns in the **Credential Profiles** window.
	Click the **Clear Filter** link to clear any filter criteria you may have set.

Create Credential Profiles

Follow the steps below to create a new credential profile. You can then use the profile to apply credentials consistently when you add new devices or providers. You can add as many protocols and corresponding credentials to the profile as you want.

If you have many credential profiles to add, you may find it more efficient to put the information in a CSV file and import the file. See Import Credential Profiles, on page 27.

When creating device credential profiles that contain SNMP credentials, Cisco recommends that the profile contain credentials for the version of SNMP actually enabled on the device, and that version only. For example: If SNMPv3 is not enabled in the device configuration, do not include SNMPv3 credentials in the device credential profile.

If you plan to use the import and export features and CSV files to create credential profiles in bulk, please note that:

- All the characters in each password or community string entry in every credential profile exported to a CSV file are replaced with asterisks (Export Credential Profiles, on page 30).
- You cannot import credential profiles if the passwords and community strings in the CSV file are blank (see Import Credential Profiles, on page 27).

To maintain network security, Cisco recommends that you use asterisks in place of real passwords and community strings in any CSV file you plan to import. After the import, follow the steps in Edit Credential Profiles, on page 30 to replace the asterisks with actual passwords and community strings.

Step 1 From the main menu, choose **Inventory Management** > **Credential Profiles** > **Credentials**.

Step 2 Click .

Step 3 In the **Profile Name** field, enter a descriptive profile name. The name can contain a maximum of 128 alphanumeric characters, plus underscores ("_") or hyphens ("-"). No other special characters are allowed.

If you will have many credential profiles, make the name as informative as possible because that information will be displayed on the Credential Profiles panel.

Step 4 Select a protocol from the **Connectivity Type** dropdown.

Step 5 Complete the credentials fields described in the following table. The required and optional fields displayed will vary with the connectivity type you chose. The values you enter must match the values configured on the device.

Connectivity Type	Fields
SSH	Enter the required **User Name**, **Password**, and **Confirm Password**. The **Enable Password** is optional.
SNMPv2	Enter the required SNMPv2 **Read Community** string. The **Write Community** string is optional.
NETCONF	Enter the required **User Name**, **Password**, and **Confirm Password**.
TELNET	Enter the required **User Name**, **Password**, and **Confirm Password**. The **Enable Password** is optional.
HTTP	Enter the required **User Name**, **Password**, and **Confirm Password**.
HTTPS	Enter the required **User Name**, **Password**, and **Confirm Password**.
GRPC	Enter the required **User Name**, **Password**, and **Confirm Password**.
SNMPv3	Choose the required **Security Level** and enter the **User Name**. If you chose the NO_AUTH_NO_PRIV **Security Level** of AUTH_NO_PRIV or AUTH_PRIV, the remaining fields are optional. If you chose the AUTH_NO_PRIV **Security Level**, you must choose an **Auth Type** and enter an **Auth Password**. If you chose the AUTH_PRIV **Security Level**, you must choose an **Auth Type** and **Priv Type**, and enter an **Auth Password** and **Priv Password**. Only the following SNMPv3 Privacy Types are supported • CFB_AES_128 • CBC_DES_56 The following Privacy Types are not supported: • AES192 • AES256 • 3DES

Step 6 (Optional) Click **+ Add Another** and repeat the above steps, as needed, for all other protocols and corresponding credentials you want to add to this credential profile.

Step 7 Click **Save**.

Import Credential Profiles

Complete the steps below to create a CSV file that specifies multiple credential profiles and then import it into Cisco Crosswork Optimization Engine.

Importing credential profiles from a CSV file adds any profiles not already in the database. You cannot import a credential profile that already exists.

If you are re-importing a credential profile CSV file that you previously exported and modified, remember that all the passwords and community strings in the exported credential profile CSV file are replaced with asterisks. You cannot re-import an exported credential profile CSV file with blank passwords. To maintain security, Cisco recommends that you use asterisks in place of real passwords and community strings in the CSV file. After the import, follow the steps in Edit Credential Profiles, on page 30 to replace the asterisks with actual passwords and community strings.

Step 1 From the main menu, choose **Inventory Management** > **Credentials**.

Step 2 Click to open the **Import CSV File** dialog box.

Step 3 If you have not already created a credential profile CSV file to import:

a) Click the **Download sample 'Credential template (*.csv)' file** link and save the CSV file template to your local disk.

b) Open the template using your preferred tool. Begin adding rows to the file, one row for each credential profile.

Use a semicolon to separate multiple entries in the same field. Use two semicolons with no space between them to indicate that you are leaving the field blank. When you separate multiple entries with semicolons, remember that the order in which you enter values in each field is important. For example, if you enter **`SSH;NETCONF;TELNET`** in the **Connectivity Type** field and you enter **`UserTom;UserDick;UserHarry;`** in the **User Name** field, the order of entry determines the mapping between the two fields:

- SSH: UserTom
- NETCONF: UserDick
- TELNET: UserHarry

Also note:

- Be sure to enter SNMP community string information exactly as currently entered on your devices. Failure to do so may result in loss of device connectivity.
- Password and community string information associated with a user ID are stored in plain text in the CSV file you prepare. Be aware of the security implications of this, and apply appropriate safeguards.

Field	Entries	Required or Optional
Credential Profile	The name of the credential profile. For example: **`srpce`**.	Required
Connectivity Type	Valid values are: **`SSH`**, **`SNMPv2`**, **`NETCONF`**, **`TELNET`**, **`HTTP`**, **`HTTPS`**, **`GRPC`** or **`SNMPv3`**	• Devices—SNMP and SSH (to avoid operational errors due to clock synchronization checks) are required. • SR-PCE—Since SR-PCE is considered a provider and a device, SSH, and HTTP are required.
User Name	For example: **`SRPCEUser`**	Required if **Connectivity Type** is **`SSH`**, **`NETCONF`**, **`TELNET`**, **`HTTP`**, **`HTTPS`**, **`SNMPv3`** or **`GRPC`**.

Field	Entries	Required or Optional
Password	The password for the preceding **User Name**.	Required if **Connectivity Type** is `SSH`, `NETCONF`, `TELNET`, `HTTP`, `HTTPS` or `GRPC`
Enable Password	Use an Enable password. Valid values are: `ENABLE`, `DISABLE`, or leave blank (unselected)	
Enable Password Value	Specify the Enable password to use.	Required only if **Enable Password** is set to `Enable`.
SnmpV2 Read Community	For example: `readprivate`	Required if **Connectivity Type** is `SNMPv2`
SnmpV2 Write Community	For example: `writeprivate`	
SnmpV3 User Name	For example: `DemoUser`	Required if **Connectivity Type** is `SNMPv3`
SnmpV3 Security Level	Valid values are `noAuthNoPriv`, `AuthNoPriv` or `AuthPriv`	Required if **Connectivity Type** is `SNMPv3`
SnmpV3 Auth Type	Valid values are `HMAC_MD5` or `HMAC_SHA`	Required if **Connectivity Type** is `SNMPv3` and **SnmpV3 Security Level** is `AuthNoPriv` or `AuthPriv`
SnmpV3 Auth Password	The password for this authorization type.	Required if **Connectivity Type** is `SNMPv3` and **SnmpV3 Security Level** is `AuthNoPriv` or `AuthPriv`
SnmpV3 Priv Type	Valid values are `CFB_AES_128` or `CBC_DES_56` The following SNMPv3 privacy types are not supported: AES192, AES256, 3DES	Required if **Connectivity Type** is `SNMPv3` and **SnmpV3 Security Level** is `AuthPriv`
SnmpV3 Priv Password	The password for this privilege type.	Required if **Connectivity Type** is `SNMPv3` and **SnmpV3 Security Level** is `AuthPriv`

Be sure to delete the sample data rows before saving the file or they will be imported along with the data you want. The column header row can stay, as it is ignored during import.

c) When you are finished, save the new CSV file.

Step 4 Click **Browse** to navigate to the CSV file you just created and then click **Open** to select it.

Step 5 With the CSV file selected, click **Import**.

The credential profiles you imported should now be displayed in the **Devices** window.

Edit Credential Profiles

A credential profile can be shared by multiple devices, even hundreds of devices in a large network. Complete the following procedure to edit credential profile settings.

Warning Changing the settings in a credential profile without first changing the settings on the device associated with the profile may result in a loss of connectivity.

Before editing any credential profile, it is always good practice to export a CSV backup of the profiles you want to change (see Export Credential Profiles, on page 30).

Step 1 From the main menu, choose **Inventory Management** > **Credentials**.

Step 2 From the left-hand side of the **Credential Profiles** window, select the profile you want to update, and click .
The **Edit Profile** window of the selected credential is displayed.

Step 3 Make the necessary changes and then click **Save**.

Delete Credential Profiles

Follow the steps below to delete a credential profile.

Note You cannot delete a credential profile that is associated with one or more devices or providers.

Step 1 Export a backup CSV file containing the credential profile you plan to delete (see Export Credential Profiles, on page 30).

Step 2 Check whether any devices or providers are using the credential profile you plan to delete. You can do this by filtering on the **Credential Profile** column, which is available on both the **Devices** window (choose **Inventory Management** > **Credentials**) and the Providers window (choose **Inventory Management** > **Credentials**).

Step 3 Reassign the devices or providers to a different credential profile (for help with this task, see Change a Device's Credential Profile, on page 31 or Change the Credential Profile for Multiple Devices, on page 31, and Edit Providers, on page 46).

Step 4 After all devices and providers have had their credential profiles reassigned: From the main menu, choose **Inventory Management** > **Credentials**.

Step 5 In the **Credential Profiles** window, choose the profile that you want to delete and then click .

Export Credential Profiles

Exporting credential profiles stores all the profiles you selected in a CSV file. This is a quick way to make backup copies of your credential profiles. You can also edit the CSV file as needed, and re-import it to add new credential profile data. You cannot overwrite existing credential profiles by importing a CSV file.

The exported credential profiles CSV file does not contain real passwords or community strings. All the characters in the passwords and community strings entries in the credential profiles are replaced with asterisks in the exported CSV file. If you plan on modifying your exported CSV file and then re-importing it, Cisco recommends that you use asterisks in place of real passwords and community strings. After the import, follow the steps in Edit Credential Profiles, on page 30 to replace the asterisks with actual passwords and community strings.

Step 1 From the main menu, choose **Inventory Management** > **Credentials**.

Step 2 (Optional) In the **Credential Profiles** window, filter the credential profile list as needed.

Step 3 Check the check boxes for the profiles you want to export. Check the check box at the top of the column to select all the profiles for export.

Step 4 Click . Depending on your browser, you will be prompted to select a path and file name to use when saving the CSV file, or to open it immediately

Change a Device's Credential Profile

You can edit device information, including changing the credential profile in the device record. This operation changes an existing association between a device and a credential profile.

Before you begin

You need a credential profile to complete this task. To create a credential profile, see Create Credential Profiles, on page 26.

Note Make sure the profile's credential settings are correct before following this procedure.

Step 1 From the main menu, choose **Device Management** > **Devices**.

Step 2 (Optional) Filter the device list by entering text in the **Search** field or filtering specific columns.

Step 3 Check the check box of the device you want to change, and click .

Step 4 Choose a different credential profile from the **Credential Profile** drop-down list.

Step 5 Click **Save**.

After the device record is updated, the system attempts to communicate with the device using the new profile. Confirm that the device is reachable without any errors.

Change the Credential Profile for Multiple Devices

If you want to change the credential profile for a large number of network devices, you may find it more efficient to make the change by editing a devices CSV file. The basic method is:

1. Export a CSV file containing the devices whose credential profiles you want to change (see Export Devices, on page 57).
2. Edit the CSV file, changing the credential profile for each device (this credential profile must already exist). Save the edited file.
3. Import the edited devices CSV file using the **Update Existing** option. You will overwrite the credential profile data for each device (see Import Devices, on page 49).

You will need to make sure that the credential profile to which you are changing already exists. If you have not yet created that credential profile, the CSV import will fail. The credential profile you associate with these devices must also have the authorization credentials for every protocol that was configured for these devices during onboarding. If any credential for a specific protocol configured on the devices is missing from or incorrect in the credential profile, then the CSV import will succeed, but reachability checks will fail for these devices.

Step 1 From the main menu, choose **Inventory Management** > **Devices**.

Step 2 Choose the devices whose credential profiles you want to change. Your options are:

- Click to include all devices.
- Filter the device list by entering text in the **Search** field or by filtering specific columns. Then click to include only the filtered list of devices.
- Check the boxes next to the device records you want to change. Then click to include only the devices that have been checked.

Step 3 Edit and save the new CSV file using the tool of your choice. Be sure to enter the correct credential profile name in the **Credential Profile** field for each device.

Step 4 Click .

Step 5 In the **Import** dialog box, click **Browse**, choose the new CSV file, and click **Update Existing**.

Manage Providers

Cisco Crosswork Optimization Engine communicates with SR-PCE and NSO providers. Cisco Crosswork Optimization Engine stores the provider connectivity details and makes that information available to applications.

Note Other providers are available on the UI. However, they are not used by Cisco Crosswork Optimization Engine. They are used by other Cisco Network Automation applications.

From the **Providers** window, you can add a new provider, update the settings configured for an existing provider, and delete a particular provider. To open this window, choose **Inventory Management** > **Providers** .

Figure 4: Providers window

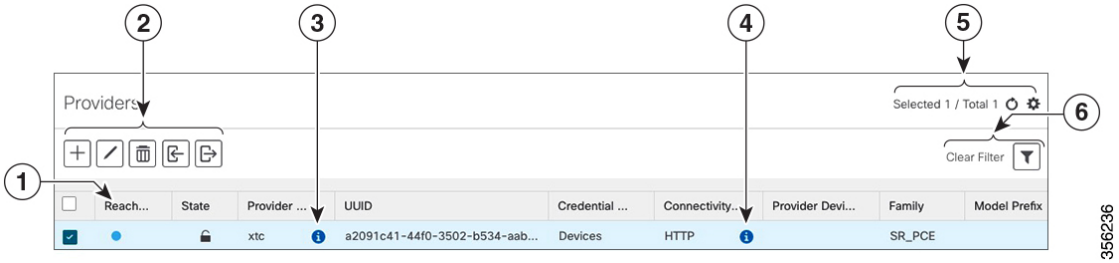

Item	Description
1	The icon shown next to the provider in this column indicates the provider's **Reachability**. For more on the icons and how reachability is determined, see Reachability and Operational State, on page 23.
2	Click to add a provider. See Add Cisco SR-PCE Providers, on page 34.
	Click to edit the settings for the selected provider. See Edit Providers, on page 46.
	Click to delete the selected provider. See Delete Providers, on page 46.
	Click to import new providers or update existing providers from a CSV file. You can also download a CSV file template by clicking this icon. The template includes sample data that you can use as a guide for building your own CSV file. See Import Providers, on page 43.
	Click to export a provider to a CSV file. See Export Providers, on page 47.
3	Click next to the provider in the **Provider Name** column to open the **Properties for** pop-up window, showing the details of any startup session key/value pairs for the provider.
4	Click next to the provider in the **Connectivity Type** column to open the **Connectivity Details** pop-up window, showing the protocol, IP and other connection information for the provider.
5	Click to refresh the **Providers** window.
	Click to choose the columns to make visible in the Providers window (see Set, Sort and Filter Table Data, on page 7).
6	Click to set filter criteria on one or more columns in the **Providers** window.
	Click the **Clear Filter** link to clear any filter criteria you may have set.

Add Cisco SR-PCE Providers

Cisco Segment Routing Path Computation Elements (Cisco SR-PCE) providers supply device discovery, management, configuration-maintenance and route-calculation services to Cisco Crosswork Optimization Engine. At least one SR-PCE provider is required in order to learn and discover SR policies, Layer 3 links, and devices.

Follow the steps below to add (through the UI) up to two instances of Cisco SR-PCE as providers for Cisco Crosswork Optimization Engine.

Before you begin

You will need to:

- Create a credential profile for the Cisco SR-PCE provider (see Create Credential Profiles, on page 26). This should be a basic HTTP text-authentication credential (currently, MD5 authentication is not supported). If the Cisco SR-PCE server you are adding does not require authentication, you must still supply a credential profile for the provider, but it can be any profile that does not use the HTTP protocol.
- Know the name you want to assign to the Cisco SR-PCE provider. This is usually the DNS hostname of the Cisco SR-PCE server.
- Know the Cisco SR-PCE server IP address.
- Determine whether you want to auto-onboard the devices that Cisco SR-PCE discovers and, if so, whether you want the new devices to have their management status set to **`managed`** or **`unmanaged`** when added. For more information, see Auto-Onboard Property Descriptions, on page 36.
- If you plan to auto-onboard devices that the Cisco SR-PCE provider discovers, and set them to a managed state when they are added to the database:
  - Assign an existing credential profile for communication with the new managed devices.
  - The credential profile must be configured with an SNMP protocol.
- For high availability, ensure that you set up two separate Cisco SR-PCE providers with unique names and IP addresses, but with matching configurations (see Configure Redundant Cisco SR-PCEs, on page 37).

Step 1 From the main menu, choose **Inventory Management** > **Providers** .

Step 2 Click [+].

Step 3 Enter the following values for the Cisco SR-PCE provider fields:

a) Required fields:

- **Provider Name**: Name of the SR-PCE provider that will be used in Cisco Crosswork Optimization Engine.
- **Credential Profile**: Select the previously created Cisco SR-PCE credential profile.
- **Family**: Select **`SR_PCE`**. All other options should be ignored.
- **Protocol**: Select **`HTTP`**. All other options should be ignored.
- **IP Address/ Subnet Mask**: Enter the IP address (IPv4 or IPv6) and subnet mask of the server.

- **Provider Properties**: Enter one of the following key/value pairs in the first set of fields (see About Adding Devices, on page 19 and Auto-Onboard Property Descriptions, on page 36):

Property Key	Value
`auto-onboard`	`off`
`auto-onboard`	`unmanaged`
`auto-onboard`	`managed`

If you enter the **`auto-onboard`**/**`managed`** pair:

1. Click the ⊕ next to the first set of fields to add a new set.
2. In the new **Property Key** field, enter **`device-profile`**.
3. In the new **Property Value** field, enter the name of a credential profile that contains SNMP credentials for all the new devices.

b) Optional value:

- **Timeout**: The amount of time (in seconds) to wait before timing out the connection to the SR-PCE server. The default is 30 seconds.

Step 4 When you have completed entries in all of the required fields, click **Save** to add the SR-PCE provider.

Step 5 Confirm that the SR-PCE provider shows a green Reachability status without any errors. You can also view the Events window to see if the provider has been configured correctly.

Note It is not recommended to modify auto-onboard options (**`managed`**/**`unmanaged`**/**`off`**) once set. If you need to modify them, do the following:

1. Delete the provider and wait until deletion confirmation is displayed in the Events page.
2. Re-add the provider with the updated auto-onboard option.
3. Confirm the provider has been added with the correct auto-onboard option in the Events page.

What to do next

- If you entered the **`auto-onboard`**/**`off`** pair, navigate to **Inventory Management** > **Devices** to add a device list (see Import Devices, on page 49).
- If you opted to automatically onboard devices, navigate to **Inventory Management** > **Devices** to view the device list. To add more node information such as geographical location details, export the device list (.csv), update it, and import it back. If geographical location data is missing, you will only be able to see device topology using the logical map.

Auto-Onboard Property Descriptions

The following table describes auto-onboard property provider fields.

Field	Description
off	If this option is enabled, you add or import devices manually (typically using a .csv file). When devices are discovered, the device data is recorded in the Cisco SR-PCE database, but is not registered in Crosswork Optimization Engine Inventory Management database.
unmanaged	If this option is enabled, all devices that Cisco SR-PCE discovers will be registered in the Cisco Crosswork Optimization Engine Inventory Management database, with their configured state set to **`unmanaged`**. SNMP polling will be disabled for these devices, and no management IP information will be included. To get these devices into the **`managed`** state later, you will need to download them as a CSV file (see Export Devices, on page 57), and modify the CSV file to add the SNMP and management IP address information. You can then update the auto-onboarded devices with this information by importing the modified CSV file (see Import Devices, on page 49). You can also assign credential profile by adding them to the device CSV file before import (the credential profiles must already exist).
managed	If this option is enabled, all devices that Cisco SR-PCE discovers will be registered in the Cisco Crosswork Optimization Engine Inventory Management database, with their configured state set to **`managed`**. SNMP polling will be enabled for these devices, and Cisco SR-PCE will also report the management IP address (Router ID). You will also need to add a second **Provider Properties** key/value pair, with the key **`device-profile`** and the value being the name of a credential profile for the new devices.

Note If **managed** or **unmanaged** options are set and you want to delete a device later, you must do one of the following:.

- Reconfigure and remove the devices from the network before deleting the device from Cisco Crosswork Optimization Engine. This avoids Cisco Crosswork Optimization Engine from rediscovering and adding the device back to Cisco Crosswork Optimization Engine.
- Set auto-onboard to **off**, and then delete the device from Cisco Crosswork Optimization Engine. However, doing so will not allow Cisco Crosswork Optimization Engine to detect or auto-onboard any new devices in the network.

Cisco SR-PCE Reachability Issues

You can find reachability issues raised in the Events table and reachability status in the **Providers** window (see Get Provider Details, on page 45). If the SR-PCE goes down, all links in the topology will display with the last known state since the SR-PCE cannot send any notification updates. When the SR-PCE becomes reachable again, a message will show in the **Events** window that SR-PCE is reconnected and the topology will be updated accordingly. If you find that the SR-PCE goes down for an extended amount of time, it is not syncing, updates are not happening, then delete the SR-PCE and add it back (when connectivity returns) using the UI:

1. Execute the following command:

```
# process restart pce_server
```

2. From the UI, navigate to **Inventory Management** > **Providers** and delete the SR-PCE provider and then add it back again.

You can also troubleshoot reachability as follows:

Step 1 Check device credentials.

Step 2 Ping the provider host.

Step 3 Attempt a connection using the protocols specified in the connectivity settings for the provider. For an SR-PCE provider, it is typically HTTP and port 8080.

```
curl --raw -vN "http://<hostname or ip-address>:8080/topology/subscribe/txt"
curl --raw -vN "http://<username>:<password>@
```

Step 4 Check your firewall setting and network configuration.

Step 5 Check the Cisco SR-PCE host or intervening devices for Access Control List settings that might limit who can connect.

Configure Redundant Cisco SR-PCEs

You can set up two Cisco SR-PCEs to ensure high availability (HA). The two Cisco SR-PCE providers must have matching configurations, supporting the same network topology. In HA, if the primary SR-PCE becomes unreachable, Cisco Crosswork Optimization Engine uses the secondary SR-PCE to discover the network topology. The network topology will continue to be updated correctly and you can view SR-PCE connectivity events in the Events table.

Configure HA

The following configurations must be done to enable HA when two Cisco SR-PCE providers are added in Cisco Crosswork Optimization Engine.

Note There must be resilient IPv4 connectivity between both SR-PCEs to enable HA. PCE IP address of the other SR-PCE should be reachable by the peer at all times.

Issue the following commands on *each* of the Cisco SR-PCE devices:

Enable the interface:

```
# interface <interface><slot>/<port>
ipv4 address <sync-link-interface-ip-address> <subnet-mask>
no shut
```

Enable HA:

```
# pce rest sibling ipv4 <other-node-pce-address>
```

Establish a sync link between the two SR-PCEs:

```
# router static
address-family ipv4 unicast
<other-node-pce-ip-address>/<subnet-mask-length> <remote-sync-link-ip-address>
```

(Optional) `# pce segment-routing traffic-eng peer ipv4 <other-node-pce-ip-address>`

It should be entered for each PCC and not for other PCE nodes.

Issue the following command on the PCC:

For SR Policies: `# segment-routing traffic-eng pcc redundancy pcc-centric`

For RSVP-TE Tunnels: `# mpls traffic-eng pce stateful-client redundancy pcc-centric`

Confirm Sibling SR-PCE Configuration

From the SR-PCE, enter the `show tcp brief` command to verify synchronization between SR-PCEs in HA are intact:

```
#show tcp brief | include <remote-SR-PCE-router-id>
```

Confirm that following information is correct:

Local Address	Foreign Address	State
<local-SR-PCE-router-id>:8080	*<local-SR-PCE-router-id>*:*<any-port-id>*	ESTAB
<local-SR-PCE-router-id>:*<any-port-id>*	*<local-SR-PCE-router-id>*:8080	ESTAB

SR-PCE Delegation

Depending on where an SR policy is created, the following SR-PCE delegation occurs:

- SR-PCE initiated—Policies configured on a PCE. SR policies are delegated back to the source SR-PCE.

  Note
  - The policy can be PCE initiated even if it is created using the UI, but in that case it is not configured explicitly on SR-PCE.
  - RSVP-TE tunnels cannot be configured directly on a PCE.

- PCC initiated—An SR policy or RSVP-TE tunnel that is configured directly on a device. The SR-PCE configured with the lowest precedence is the delegated SR-PCE. If precedence is not set, then SR-PCE with the lowest PCE IP address is the delegated SR-PCE. The following configuration example, shows that **`10.0.0.1`** is assigned a precedence value of 10 and will be the delegated SR-PCE.

```
segment-routing
  traffic-eng
    pcc
      source-address ipv4 10.0.0.2
      pce address ipv4 10.0.0.1
        precedence 10
       !
      pce address ipv4 10.0.0.8
        precedence 20
       !
       report-all
       redundancy pcc-centric
```

  For RSVP-TE Tunnel:

```
mpls traffic-eng
interface GigabitEthernet0/0/0/0
  admin-weight 1
!
```

```
interface GigabitEthernet0/0/0/1
  admin-weight 1
!
interface GigabitEthernet0/0/0/2
  admin-weight 1
!
pce
  peer source ipv4 192.168.0.02
  peer ipv4 192.168.0.9
    precedence 10
  !
  peer ipv4 192.168.0.10
    precedence 20
  !
  stateful-client
   instantiation
   report
   redundancy pcc-centric
   autoroute-announce
  !
!
auto-tunnel pcc
  tunnel-id min 990 max 999
```

- Cisco Crosswork Optimization Engine SR-PCE initiated—An SR policy that is configured using Cisco Crosswork Optimization Engine. SR-PCE delegation is random per policy.

Note Only TE tunnels (SR policies or RSVP-TE tunnels) created by Cisco Crosswork Optimization Engine can be modified or deleted by Cisco Crosswork Optimization Engine.

HA Notes and Limitations

- It is assumed that all PCCs are PCEP connected to both SR-PCEs.
- When an SR-PCE is disconnected only from Cisco Crosswork Optimization Engine, the following occur:
  - SR-PCE delegation assignments remain, but the SR-PCE that has been disconnected will not appear in Cisco Crosswork Optimization Engine.
  - You are not able to modify Cisco Crosswork Optimization Engine SR-PCE initiated SR policies if the disconnected SR-PCE is the delegated PCE.
- After an SR-PCE reloads, do the following:
  1. Execute the following command:

     ```
     # process restart pce_server
     ```

  2. From the UI, navigate to **Inventory Management** > **Providers** and delete the PCE sibling configuration in both SR-PCEs and then add the sibling configuration back again.
- In some cases, when an SR policy that was created via the UI is automatically deleted (intentional and expected) from Cisco Crosswork Optimization Engine, a warning message does not appear. For example, if the source PCC is reloaded, the UI created SR policy disappears and the user is not informed.

- In an extreme case where one SR-PCE fails on all links (to PCCs/topology devices) except the up-link to Cisco Crosswork Optimization Engine, then topology information will not be accurate in Cisco Crosswork Optimization Engine. When this happens, fix the connectivity issue or delete both SR-PCEs from the Provider page and re-add the one that is reachable.

SR-PCE Configuration Examples

The following configurations are *examples* to guide you in a multiple SR-PCE setup for HA. Please modify accordingly.

Sample redundant SR-PCE configuration (on PCE)

```
pce
 address ipv4 192.168.0.7
 rest
  sibling ipv4 192.168.0.6
```

Sample redundant SR-PCE Configuration (PCC)

```
segment-routing
 traffic-eng
  pcc
   source-address ipv4 192.0.2.1
   pce address ipv4 192.0.2.6
    precedence 200
   !
   pce address ipv4 192.0.2.7
    precedence 100
   !
   report-all
   redundancy pcc-centric
```

Sample redundant SR-PCE Configuration (on PCC) for RSVP-TE

Note `Loopback0` represents the TE router ID.

```
ipv4 unnumbered mpls traffic-eng Loopback0
!
mpls traffic-eng
 pce
  peer source ipv4 209.165.255.1
  peer ipv4 209.165.0.6
   precedence 200
  !
  peer ipv4 209.165.0.7
   precedence 100
  !
  stateful-client
   instantiation
   report
   redundancy pcc-centric
   autoroute-announce
  !
 !
 auto-tunnel pcc
  tunnel-id min 1000 max 1999
 !
!
```

Sample SR-TM Configuation

```
telemetry model-driven
 destination-group crosswork
  address-family ipv4 198.18.1.219 port 9010
   encoding self-describing-gpb
   protocol tcp
  !
 !
 sensor-group SRTM
  sensor-path Cisco-IOS-XR-infra-tc-oper:traffic-collector/afs/af/counters/tunnels
  sensor-path
Cisco-IOS-XR-infra-tc-oper:traffic-collector/vrf-table/default-vrf/afs/af/counters/prefixes

 !
 subscription OE
  sensor-group-id SRTM sample-interval 60000
  destination-id crosswork
  source-interface Loopback0
!
traffic-collector
 interface GigabitEthernet0/0/0/3
 !
 statistics
  history-size 10
```

Note The destination address uses the southbound data interface (eth1) address of the Cisco Crosswork Data Gateway VM.

Path Computation Client (PCC) Support

PCCs can support delegation and reporting of both RSVP-TE tunnels and SR policies to SR-PCE. In order for both to be supported on the same PCC, two separate PCEP connections must be established with the SR-PCEs. Each PCEP connection must have a distinct source IP address (Loopback) on the PCC.

The following is a Cisco IOS-XR configuration example of PCEP connections for RSVP-TE, where 192.168.0.2 is the PCEP session source IP for RSVP-TE tunnels delegated and reported to SR-PCE. It is a loopback address on the router. Two SR-PCEs are configured for PCEP sessions, where the first will be preferred for delegation of RSVP-TE tunnels due to precedence. Auto-tunnel PCC is configured with a range of tunnel IDs that will be used for assignment to PCE-initiated RSVP-TE tunnels like those created in Crosswork Optimization Engine.

```
mpls traffic-eng
interface GigabitEthernet0/0/0/2
admin-weight 1
!
interface GigabitEthernet0/0/0/3
admin-weight 1
  pce
    peer source ipv4 192.168.0.2
    peer ipv4 192.168.0.1
      precedence 10
     !
    peer ipv4 192.168.0.8
      precedence 11
     !
```

```
    stateful-client
      instantiation
      report
     !
   !
   auto-tunnel pcc
    tunnel-id min 10 max 1000
   !
!
ipv4 unnumbered mpls traffic-eng Loopback0

rsvp
interface GigabitEthernet0/0/0/2
bandwidth 1000000
!
interface GigabitEthernet0/0/0/3
bandwidth 1000000
!
!
```

Add Cisco NSO Providers

Cisco Network Services Orchestrator (Cisco NSO) providers supply device management and configuration maintenance services to Cisco Crosswork Optimization Engine.

Follow the steps below to add (through the UI) a Cisco NSO provider for Cisco Crosswork Optimization Engine. You can also add providers using CSV files (see Import Providers, on page 43).

Before you begin

You will need to:

- Create a credential profile for the Cisco NSO provider (see Create Credential Profiles, on page 26).

  Know the name you want to assign to the Cisco NSO provider.
- Know the Cisco NSO NED device models and driver versions used in your topology.
- Know the Cisco NSO server IP address and hostname.
- Confirm Cisco NSO device configurations (see Sample Configuration for Devices in Cisco NSO, on page 22).

Step 1 From the main menu, choose **Inventory Management** > **Providers** .

Step 2 Click [+].

Step 3 Enter the following values for the Cisco NSO provider fields:

a) Required fields:

- **Provider Name**: Enter a name for the provider that will be used in Cisco Crosswork Optimization Engine.
- **Credential Profile**: Select the previously created Cisco NSO credential profile.
- **Family**: Select `NSO`.

- **Device Key**: Select the method that Cisco NSO uses to identify devices uniquely. This will serve as the way Cisco Crosswork Optimization Engine maps the device to Cisco NSO. Choose **`NODE_IP`** and other options you wish.
- Under Connection Type(s), **Protocol**:
- **IP Address/Subnet Mask**: Enter the IP address (IPv4 or IPv6) and subnet mask of the Cisco NSO server.
- **Port**: Enter the port to use to connect to the Cisco NSO server. The default is **`2022`**.
- **Model**: Select the model (**Cisco-IOS-XR**, **Cisco-NX-OS**, or **Cisco-IOS-XE**) from the drop-down list and enter its associated NED driver version. Add a model for each type of device that will be used in the topology. If you have more than one, select + to add another supported model.
- **Version**: Enter the default software version of the device.

For more information on fields, see Import Providers, on page 43.

b) Optional values:

- **Timeout**: The amount of time (in seconds) to wait before timing out the connection to the Cisco NSO server. The default is 30 seconds.

Step 4 Under Provider Properties, enter a **Provider Key** of **`forward`** and a **Property Value** of **`true`**.

Step 5 When you have completed entries in all of the required fields, click **Save** to add Cisco NSO as a provider.

Import Providers

Complete the steps below to create a CSV file that specifies providers and then import it into Cisco Crosswork Optimization Engine.

Importing providers from a CSV file adds any providers not already in the database, and updates any providers with the same name as an imported provider. For this reason, it is a good idea to export a backup copy of all your current providers before an import (see Export Providers, on page 47).

Step 1 From the main menu, choose **Inventory Management** > **Providers** .

Step 2 Click to open the **Import CSV File** dialog box.

Step 3 If you have not already created a provider CSV file to import:

a) Click the **Download sample 'Provider template (*.csv)' file** link and save the CSV file template to a local storage resource.

b) Open the template using your preferred tool. Begin adding rows to the file, one row for each provider.

Use a semicolon to separate multiple entries in the same field. Use two semicolons with no space between them to indicate that you are leaving the field blank. When you separate entries with semicolons, the order in which you enter values is important. For example, if you enter **`SSH;SNMP;NETCONF;TELNET`** in the **connectivity_type** field and you enter **`22;161;830;23`** in the **connectivity_port** field, the order of entry determines the mapping between the two fields:

- SSH: port 22

- SNMP: port 161
- NETCONF: port 830
- Telnet: port 23

Field	Description	Required or Optional
Provider Name	Enter the name for the provider that will be used in Crosswork Optimization Engine. For example: **`MySRPCE`**.	Required
Connectivity Type	Enter the name of the protocol that Crosswork Optimization Engine will use to connect to the provider. For example: **`ROBOT_MSVC_TRANS_HTTP`** = HTTP	Required
Connectivity IP	Enter the IP address (IPv4 or IPv6) of the provider. **Note** When using an IPv6 address, please note the following: • In the Properties column, do not set the auto-onboard property to **`auto-onboard:managed`**. • The IPv6 host is detected so the deployment mode will automatically be set to IPv6.	Required
Connectivity Port	Enter the port number to use to connect to the provider's server.	Required
Connectivity Timeout	Enter the amount of time (in seconds) to wait before the connection to the provider times out. The default is 30 seconds.	Optional
Credential Profile Name	Enter the name of the credential profile that Crosswork Optimization Engine will use to connect to the provider. This profile must already exist in the system.	Required
Provider Device Key	Enter the enum value corresponding to the key that the Cisco NSO provider uses to identify devices uniquely. This will serve as the way Crosswork Optimization Engine maps the device to the Cisco NSO provider. Valid values are: • **`ROBOT_PROVDEVKEY_HOST_NAME`**—If you are using the device hostname as the device ID within NSO, this value must match the hostname that is specified for the device in the inventory. • **`ROBOT_PROVDEVKEY_NODE_IP`**—Use this enum value if the NSO device identifier is the IP address for the Node IP value in the CSV file. • **`ROBOT_PROVDEVKEY_INVENTORY_ID`**—Use this enum value if the inventory ID is the device identifier for NSO.	This entry is only required if you are creating or updating a Cisco NSO provider. Otherwise, leave blank.
Family	Enter **`ROBOT_PROVIDER_SR_PCE`** or **`ROBOT_PROVIDER_SR_NSO`**. Do not choose other options as they are reserved for use by other Cisco Network Automation applications.	Required

Field	Description	Required or Optional
Model Prefix	If you are adding a Cisco NSO provider: Select the model prefix that matches the NED CLI used by the NSO server. Valid entries are: **`Cisco-IOS-XR`**, **`Cisco-NX-OS`**, **`Cisco-IOS-XE`**. For telemetry, only Cisco-IOS-XR is supported.	Required for Cisco NSO providers only
Model Version	If you adding a Cisco NSO provider: Enter the Cisco NSO NED driver version used on the server.	Required for Cisco NSO providers only
Properties	Enter the Cisco SR-PCE appropriate auto-onboard entries: **`auto-onboard:<auto-onboard property>;device-profile:<SR-PCE Credential Profile Name>`** For example: **`auto-onboard:managed;device-profile:cisco`** When using IPv6 connectivity, do not set **`auto-onboard:managed`**. See Add Cisco SR-PCE Providers, on page 34.	This entry is only required if you are creating or updating a Cisco SR-PCE provider. Otherwise, leave blank.

Be sure to delete the sample data rows before saving the file, or they will be imported along with the data you want. The column header row can stay, as it is ignored during import.

c) When you are finished, save the new CSV file.

Step 4 Click **Browse** to navigate to the CSV file you just created and then click **Open** to select it.

Step 5 With the CSV file selected, click **Import**.

The provider information you imported should now be displayed in the **Providers** window.

Step 6 Resolve any errors reported during the import and check provider details to confirm connection.

Get Provider Details

Use the **Providers** window to get details about your providers and to check on their reachability.

Step 1 From the main menu, choose **Inventory Management** > **Providers** .
For each provider configured in Cisco Crosswork Optimization Engine, the **Providers** window lists information such as the provider's name, universally unique identifier (UUID), associated credential profile, device key, and more, as shown in the figure below.

Figure 5: Providers Window

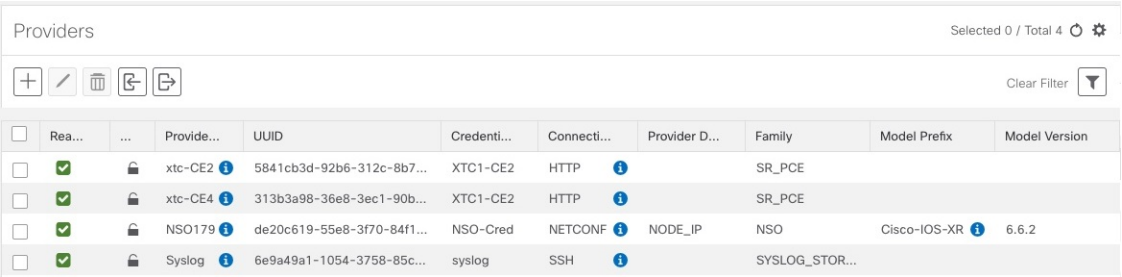

Step 2 The icons in the **Reachability** column indicate whether a provider is reachable via the listed connectivity protocols. For a description of each icon and its meaning, see Reachability and Operational State, on page 23.

Cisco Crosswork Optimization Engine checks provider reachability immediately after a provider is added or modified. Other than these events, Cisco Crosswork Optimization Engine checks SR-PCE reachability about every 10 seconds.

Step 3 Get additional details for any provider, as follows:

a) In the **Provider Name** column, click the to view provider-specific key/value properties.

b) In the **Connectivity Type** column, click the to view detailed connectivity information for the provider, such as provider-specific protocol, IP format, IP address, port, and timeout information.

c) When you are finished, click ✕ to close the details window.

If you are running into Cisco SR-PCE reachability problems, see Cisco SR-PCE Reachability Issues, on page 36.

Edit Providers

When editing provider settings, be aware that a provider can be mapped to many devices, even thousands of devices in a large network.

Note

- Before making any changes to a provider configuration you should be certain that you understand the full impact of the change. If you are unsure about the potential risk of making a change, contact Cisco services for guidance.
- See Add Cisco SR-PCE Providers, on page 34 before modifying an SR-PCE provider. There are additional steps that must be done when editing an SR-PCE provider.

.

Before editing any provider, it is always good practice to export a CSV backup of the providers you want to change (see Export Providers, on page 47).

Step 1 From the main menu, choose **Inventory Management** > **Providers** .

Step 2 In the **Providers** window, choose the provider you want to update and click .

Step 3 Make the necessary changes and then click **Save**.

Step 4 Resolve any errors and confirm provider reachability.

Delete Providers

Follow the steps below to delete a provider.

Note If an SR-PCE provider's auto-onboard **managed** or **unmanaged** options are set, you must do one of the following:.

- Reconfigure and remove the devices from the network before deleting the device from Cisco Crosswork Optimization Engine. This avoids Cisco Crosswork Optimization Engine from rediscovering and adding the device back.
- Set auto-onboard to **off**, and then delete the device from Cisco Crosswork Optimization Engine. However, doing so will not allow Cisco Crosswork Optimization Engine to detect or auto-onboard any new devices in the network.

You are alerted when you try to delete a provider that is associated with one or more devices or credential profiles.

Step 1 Export a backup CSV file containing the provider you plan to delete (see Export Providers, on page 47).

Step 2 Delete the provider as follows:

a) From the main menu, choose **Inventory Management** > **Providers** .

b) In the **Providers** window, choose the provider(s) that you want to delete and click .

c) In the confirmation dialog box, click **Delete**.

Export Providers

You can quickly export provider data to a CSV file. This is a handy way to keep backup copies of your provider information.

Note You cannot edit a CSV file and then re-import it to update existing providers.

Step 1 From the main menu, choose **Inventory Management** > **Providers** .

Step 2 (Optional) In the **Providers** window, filter the provider list as needed.

Step 3 Check the check boxes for the providers you want to export. Check the check box at the top of the column to select all the providers for export.

Step 4 Click . Depending on your browser, you will be prompted to select a path and file name to use when saving the CSV file, or to open it immediately.

View Devices Assigned to a Provider

To see a list of devices that are assigned to a particular Cisco NSO provider:

Step 1 From the main menu, choose **Inventory Management** > **Devices**.

Step 2 In the **Devices** window, scroll across the table until you find the **Providers** column.

Step 3 Under the Local Config field, enter filter criteria.
The table displays only the devices with the Provider criteria you entered.

Manage Network Devices

The Device Management application's **Network Devices** window (shown below) gives you a consolidated list of all your devices and their status. To view the **Network Devices** window, select **Inventory Management** > **Devices**. The **Network Devices** tab is displayed by default.

Figure 6: Devices Window

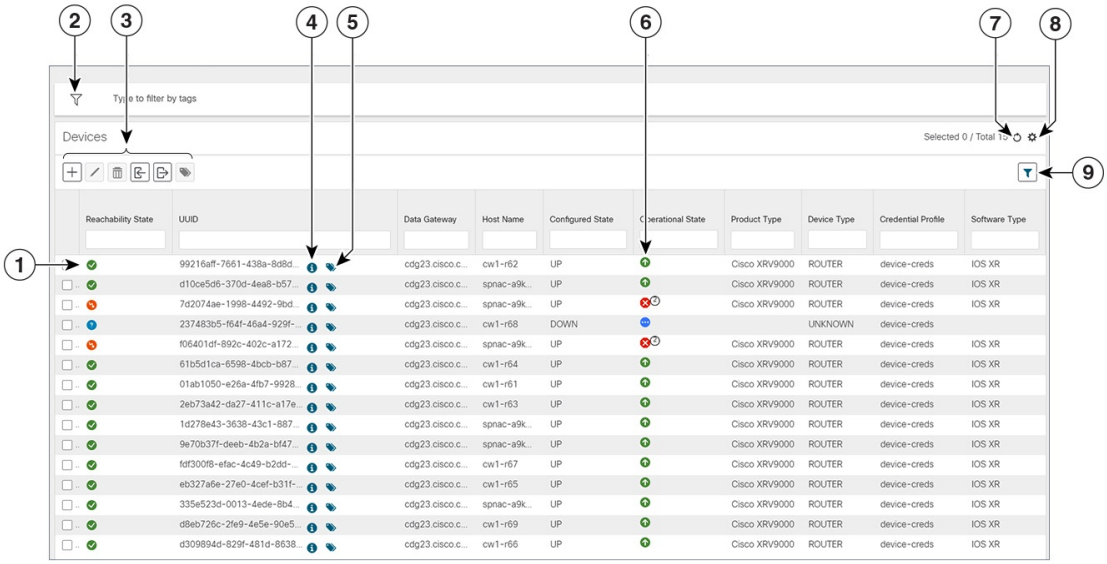

Item	Description
1	The **Filter by tags** field lets you filter the devices by the tags applied to them. Type the name of the tag that has been applied to the device that you are trying to find. See Filter Devices by Tags, on page 56.

Item	Description
2	Click to add a new device to the device inventory. See About Adding Devices, on page 19.
	Click to edit the information for the currently selected devices. See Edit Devices, on page 56.
	Click to delete the currently selected devices. See Delete Devices, on page 57.
	Click to import new devices and update existing devices, using a CSV file. You can also download a CSV file template by clicking this icon. The template includes sample data that you can use as a guide for building your own CSV file. See Import Devices, on page 49.
	Click to export information for selected devices to a CSV file. See Export Devices, on page 57.
	Click to modify tags applied to the selected devices. See Apply or Remove Device Tags, on page 62.
3	Click to open the **Device Details** pop-up window, where you can view important information for the selected device. See Get Device Details, on page 54.
4	Icons in the **Operational State** column show whether a device is operational or not. See Reachability and Operational State, on page 23
5	Click to refresh the Devices list.
6	Click to select which columns to display in the Devices list (see Set, Sort and Filter Table Data, on page 7).
7	Click to set filter criteria on one or more columns in the Devices list.
	Click the **Clear Filter** link to clear any filter criteria you may have set.
8	Icons in the **Reachability State** column show whether a device is reachable or not. See Reachability and Operational State, on page 23.

Import Devices

Complete the steps below to create a CSV file that specifies multiple devices and then import it into Cisco Crosswork Optimization Engine.

Importing devices from a CSV file adds any devices not already in the database. The **Update Existing** option overwrites the data in any device record with a device key field value that matches those of an imported device (this excludes the UUID, which is set by the system and not affected by import). For this reason, it is a good idea to export a backup copy of all your current devices before an import (see Export Devices, on page 57).

Note If you plan on using a CSV file to import devices managed by Cisco Network Services Orchestrator (Cisco NSO), you must prepare the CSV following the guidelines given in Sample Configuration for Devices in Cisco NSO, on page 22.

Step 1 From the main menu, choose **Inventory Management** > **Devices**.

Step 2 Click to open the **Import CSV File** dialog box.

Step 3 If you have not already created a device CSV file to import:

a) Click the **Download sample 'Device Management template (*.csv)' file** link and save the CSV file template to a local storage resource.

b) Open the template using your preferred tool. Begin adding rows to the file, one row for each device.

Note Confirm that the TE router ID value for each device is populated. This value is used to uniquely identify the device in the topology which is learned from SR-PCE. Without a valid TE router ID for each device, the topology will not be displayed.

Use a semicolon to separate multiple entries in the same field. Use two semicolons with no space between them to indicate that you are leaving the field blank. When you separate multiple entries with semicolons, remember that the order in which you enter values in each field is important. For example, if you enter **`SSH;SNMP;NETCONF;TELNET`** in the **Connectivity Type** field and you enter **`22;161;830;23`** in the **Connectivity Port** field, the order of entry determines the mapping between the two fields:

- SSH: port 22
- SNMP: port 161
- NETCONF: port 830
- Telnet: port 23

For a list of the fields and the mandatory values you must enter, see the "Add New Device" field table in Add Devices Through the UI, on page 51.

Be sure to delete the sample data rows before saving the file, or they will be imported along with the data you want. The column header row can stay, as it is ignored during import.

c) When you are finished, save the new CSV file.

Step 4 Click **Browse** to navigate to the CSV file you just created and then click **Open** to select it.

Step 5 With the CSV file selected, click **Import** to add new devices or **Update Existing** to add or change data to devices already in the system.

Step 6 Resolve any errors and confirm device reachability.

The device information you imported should be displayed in the **Devices** window within a few minutes (see Manage Network Devices, on page 48).

It is normal for devices to show as unreachable or not operational when they are first imported. However, if after 30 minutes they are still displayed as unreachable or not operational, there is an issue that needs to be investigated. To investigate, select **Inventory Management** > **Job History** and click on any you see in the **Status** column. Common issues include failure to ensure the associated credential profile contains the correct credentials. You can test this by

opening a terminal window on the Cisco Crosswork Optimization Engine server and then trying to access the device using the protocol and credentials specified in the associated credential profile.

Add Devices Through the UI

Follow the steps below to add devices one by one, using the UI. Under normal circumstances, you will want to use this method when adding one or a few devices only.

Before you begin

Be sure you have completed the planning steps and setup requirements discussed in Get Started, on page 11, and that the devices themselves have been pre-configured as explained in Prerequisites for Onboarding Devices, on page 21.

Step 1 From the main menu, choose **Devices Management** > **Devices**.

Step 2 Click [+].

Step 3 Enter values for the new device, as listed in the table below.

Step 4 Click **Save**. (The Save button is disabled until all mandatory fields are complete.)

Step 5 (Optional) Repeat to add more devices.

Table 4: Add New Device Window (=Required)*

Field	Description
*** Configured State**	The management state of the device. Options are • **`UNMANAGED`**—Cisco Crosswork Optimization Engine is not monitoring the device. • **`DOWN`**—The device is being managed and is down. • **`UP`**—The device is being managed and is up.
*** Reachability Check**	Determines whether Cisco Crosswork Optimization Engine performs reachability checks on the device. Options are: • **`ENABLE`** (In CSV: **`REACH_CHECK_ENABLE`**)—Checks for reachability and then updates the Reachability State in the UI automatically. • **`DISABLE`** (In CSV: **`REACH_CHECK_DISABLE`**)—The device reachability check is disabled. Cisco recommends that you always set this to **`ENABLE`**. This field is optional if **Configured State** is marked as **`UNMANAGED`**.
*** Credential Profile**	The name of the credential profile to be used to access the device for data collection and configuration changes. For example: **`nso23`** or **`srpce123`**. This field is optional if **Configured State** is marked as **`UNMANAGED`**.

Field	Description
Host Name	The hostname of the device. Cisco Crosswork Optimization Engine discovers it and updates it.
Inventory ID	Inventory ID value for the device.
UUID	Universally unique identifier (UUID) for the device.
Serial Number	Serial number for the device.
Node IP	IP address of the device.
MAC Address	MAC address of the device.
*** Capability**	The capabilities that allow collection of device data and that are configured on the device. You must select at least **SNMP** as this is a required capability. The device will not be onboarded if **SNMP** is not configured. Other options are **YANG_MDT**, **TL1**, **YANG_CLI**, and **YANG-EPNM**. The capabilities you select will depend on the device software type and version.
Tags	The available tags to assign to the device for identification and grouping purposes. Use device tags to group devices for monitoring, and to provide additional information that might be of interest to other users, such as the device's physical location or its administrator's email ID. For more information, see Manage Tags.
Connectivity Details	
Protocol	The connectivity protocols used by the device. Choices are: **SSH**, **SNMP**, **NETCONF**, **TELNET**, **HTTP**, and **HTTPS**. To add more connectivity protocols for this device, click + at the end of the first row in the **Connectivity Details** panel. To delete a protocol you have entered, click × shown next to that row in the panel. You can enter as many sets of connectivity details as you want, including multiple sets for the same protocol. You must enter details for at least **SSH** and **SNMP**. If you do not configure **SNMP**, the device will not be added. If you want to manage the device (or you are managing XR devices), you must enter details for **NETCONF**. **TELNET** connectivity is optional.
*** IP Address / Subnet Mask**	Enter the device's IP address (IPv4 or IPv6) and subnet mask.
*** Port**	The port used for this connectivity protocol. Each protocol is mapped to a port, so be sure to enter the port number that corresponds to the **Protocol** you chose. The standard port assignments for each protocol are: • SSH: 22 • SNMP: 161 • NETCONF: 830 • TELNET: 23 • HTTP: 80 • HTTPS: 443

Field	Description
Timeout	The elapsed time (in seconds) before communication attempts using this protocol will time out. The default value is 30 seconds. For XE devices using NETCONF, the recommended minimum timeout value is 90 seconds. For all other devices and protocols, the recommended minimum timeout value is 60 seconds.
Routing Info	
ISIS System ID	The device's IS-IS system ID. This ID identifies the router in an IS-IS topology, and is required for SR-PCE integration.
OSPF Router ID	The device's OSPF router ID. This ID identifies the router in an OSPF topology, and is required for SR-PCE integration.
***TE Router ID**	The device’s OSPF Router ID or ISIS Router ID depending on the IGP used in the network topology.
Streaming Telemetry Config	
Telemetry Interface Source VRF	Name of the VRF within which Model Driven Telemetry (MDT) traffic is routed.
Location All location fields are optional, with the exception of **Longitude** and **Latitude**, which are required for the geographical view of your network topology.	
Longitude, **Latitude**	Longitude and latitude values are required so that the geographical map can present the correct geographical location of the device and its links to other devices. Enter the longitude and latitude in Decimal Degrees (DD) format.
Altitude	The altitude, in feet or meters, at which the device is located. For example, `123`.
Providers and Access	
Local Config: **Device Key** and **Provider**	Mandatory only when mapping an NSO provider. The Device Key will automatically populate and the Credential Profile appears. For CSV entry, use ROBOT_PROVIDER_LOCAL_CONFIG and enter the Provider name.
Compute Config: **Provider**	(Optional) Provider name used for topology computation. Choose a provider from the list. For CSV entry, use ROBOT_PROVIDER_COMPUTE and enter the Provider name.

Example

Figure 7: Add New Device Window

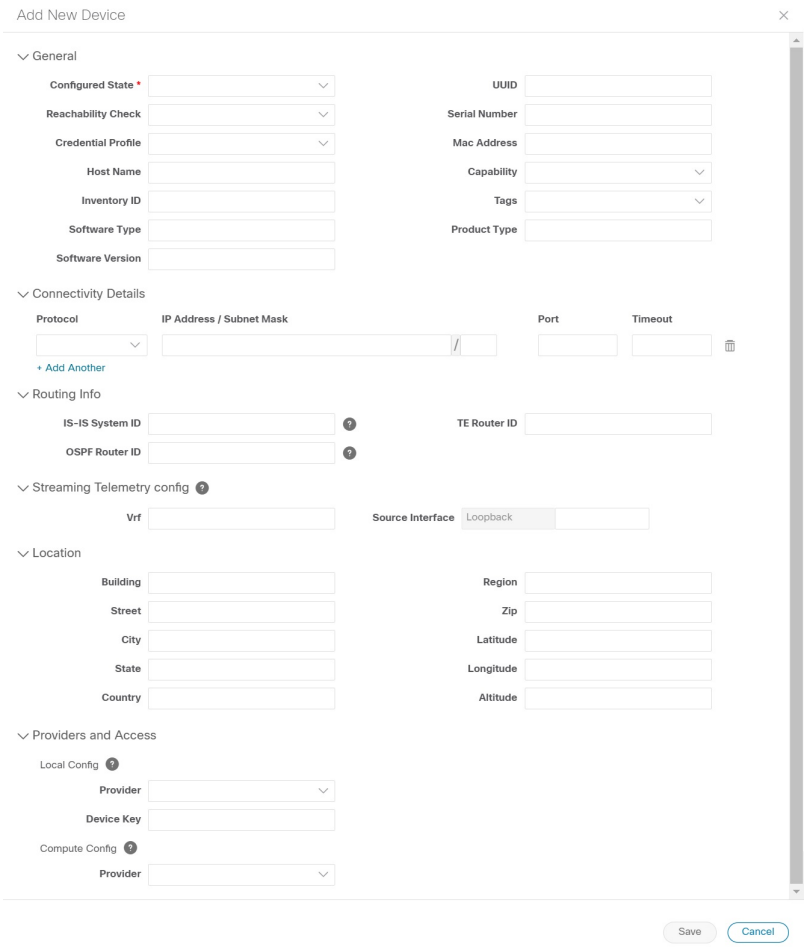

Get Device Details

Whenever you select **Inventory Management** > **Devices** and display the list of devices under the **Network Devices** tab, you can click ⓘ next to any listed device to get more information about that device. Clicking this icon opens the **Details for DeviceName** pop-up window, as shown in the following example:

Figure 8: Details for DeviceName Window

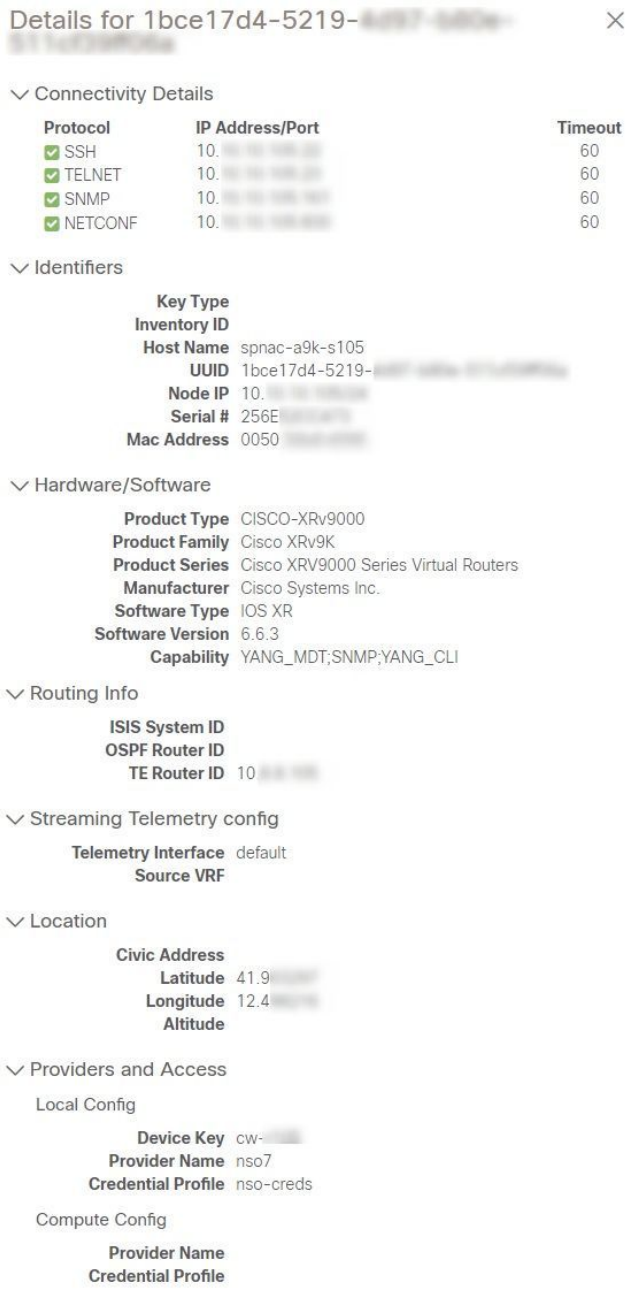

Expand the **Connectivity Details** area at the top of the pop-up window (if it is not already expanded). This area shows the reachability status for all transport types (for help with the icons shown in this area, see Device and Link Icons, on page 68).

Expand and collapse the other areas of the pop-up window, as needed. Click ✕ to close the window.

Filter Devices by Tags

By creating a tag and assigning it to a particular device, you can easily provide additional information that might be of interest to other users, such as the device's physical location and its administrator's email ID. You can also use tags to find and group devices with the same or similar tags in any window that lists devices.

For help with tagging your devices, see Apply or Remove Device Tags, on page 62. For help with creating and deleting tags, see Manage Tags, on page 59.

To filter devices by tags:

Step 1 Display the **Devices** window by choosing **Inventory Management** > **Devices**.

Step 2 In the **Type to filter by tags** bar at the top of the user interface, type all or part of the name of a tag.

The **Type to filter by Tags** bar has a type-ahead feature: As you start typing, the field shows a drop-down list of tags that match all the characters you have typed so far. To force the drop-down list to display all available tags, type *****.

Step 3 Choose the name of the tag you want to add to the filter. The filter appears in the **Type to filter by tags** filter bar. The table or map shows only the devices with that tag.

Step 4 If you want to filter on more than one tag:

a) Repeat Steps 2 and 3 for each additional tag you want to set as part of the filter.

b) When you have selected all the tags you want, click **Apply Filters**. The table or map shows only the devices with tags that match **all** the tags in your filter.

Step 5 To clear all tag filters, click the **Clear Filters** link. To remove a tag from a filter containing multiple tags, click the **X** icon next to that tag's name in the filter.

Edit Devices

Complete the following procedure to update a device's information.

Before editing any device, it is always good practice to export a CSV backup of the devices you want to change (see Export Devices, on page 57).

Step 1 From the main menu, choose **Inventory Management** > **Devices**.

Step 2 (Optional) Filter the list of devices by filtering specific columns.

Step 3 Check the check box of the device you want to change, then click .

Step 4 Edit the values configured for the device, as needed. For a description of the fields you can update, see Add Devices Through the UI.

Note In addition to the exising fields, you can also view the **Data Gateway** configured for the selected device. This field is read-only.

Step 5 Click **Save**. (The Save button remains dimmed until all required fields are filled in.)

Step 6 Resolve any errors and confirm device reachability.

Delete Devices

Complete the following procedure to delete devices.

Before you begin

- If the auto-onboard **managed** or **unmanaged** options are set for the SR-PCE provider, you should set auto-onboard for the SR-PCE(s) to **off**.
- Confirm that the device is not connected to the network or that it is powered off before deleting the device.

Note

- If devices are mapped to Cisco NSO with MDT capability, and telemetry configuration is pushed, then those configurations will be removed from the device.
- If auto-onboard is not set to **off**, and it is still functional and connected to the network, the device will be rediscovered as unmanaged as soon as it is deleted.

Step 1 Export a backup CSV file containing the devices you plan to delete (see Export Devices, on page 57).

Step 2 From the main menu, choose **Inventory Management** > **Devices**.

Step 3 (Optional) In the **Devices** window, filter the list of devices by entering text in the **Search** field or filtering specific columns.

Step 4 Check the check boxes for the devices you want to delete.

Step 5 Click to change each device's state to ADMIN DOWN or UNMANAGED.

If you want to delete devices in bulk, Cisco recommends that you change the device state in this manner in batches of 50 devices, then complete deletion of these devices before deleting another batch.

Step 6 Click .

Step 7 In the confirmation dialog box, click **Delete**.

Export Devices

When you export the device list, all device information is exported to a CSV file. Exporting the device list is a handy way to keep a record of all devices in the system at one time. You can also edit the CSV file as needed, and re-import it to overwrite existing device data.

Step 1 From the main menu, choose **Inventory Management** > **Devices**.

Step 2 (Optional) Filter the device list as needed.

Step 3 Check the check boxes for the devices you want to export. Check the check box at the top of the column to select all the devices for export.

Step 4 Click . Your browser will prompt you to select a path and the file name to use when saving the CSV file, or to open it immediately

View Device Job History

Device Management collects and stores information about device-related jobs. Follow the steps below to track all create, update and delete activities.

Step 1 From the main menu, choose **Inventory Management** > **Job History**. The **Inventory Jobs** window displays a log of all device-related jobs, like the one shown below.

Figure 9: Job History Window With Error Details Popup

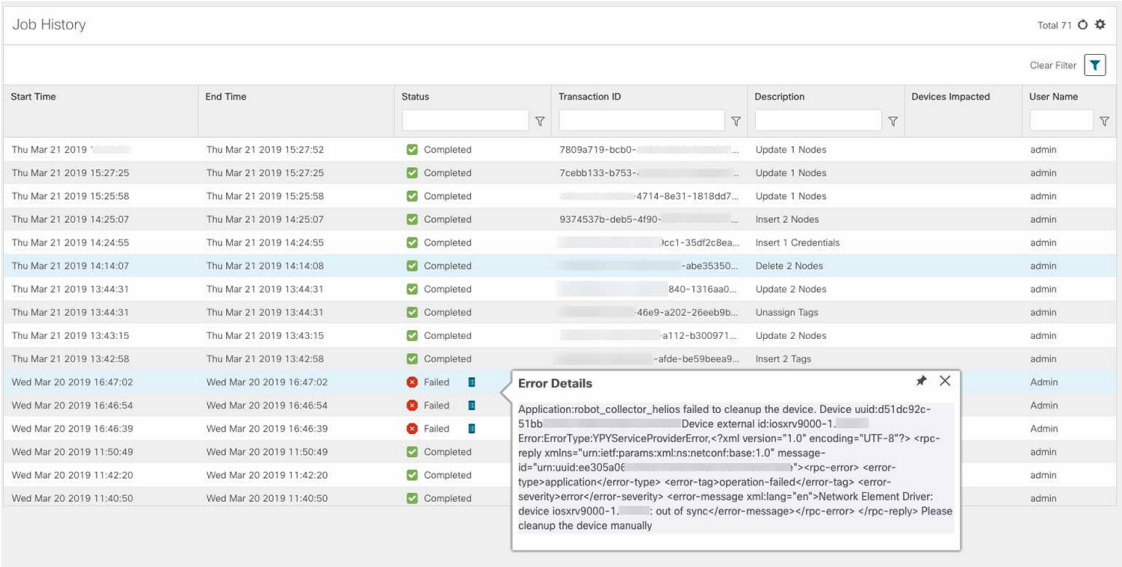

The jobs display in descending order of creation time. The most recent job is shown first. To sort the data in the table, click a column heading. You can toggle between ascending and descending sort order (for more help, see Set, Sort and Filter Table Data, on page 7).

Step 2 The **Status** column shows three types of states: completed, failed, and partial. For any failed or partial job, click shown next to the error for information.

Error information may include `clean-up failure` events as audit messages. These messages indicate that Cisco Crosswork Network Automation configuration objects on the device could not be removed, and will explain why they could not be removed. Users will need to take manual action to remove them. This typically involves deleting any XR telemetry configuration objects with names starting with `CW_`.

Manage Tags

Use the **Tag Management** window to manage the tags available for assignment to the devices in your network. Tags can provide information such as the device's physical location and its administrator's email ID, and are used to group devices.

To open this window, choose **Admin** > **Tags** from the main window.

Note Cisco Crosswork Optimization Engine automatically creates a default set of tags and assigns them to every device it manages:

- cli
- mdt
- reach-check
- snmp
- clock-drift-check

You cannot select, edit, delete, or manually associate these default tags with any device.

Figure 10: Tag Management Window

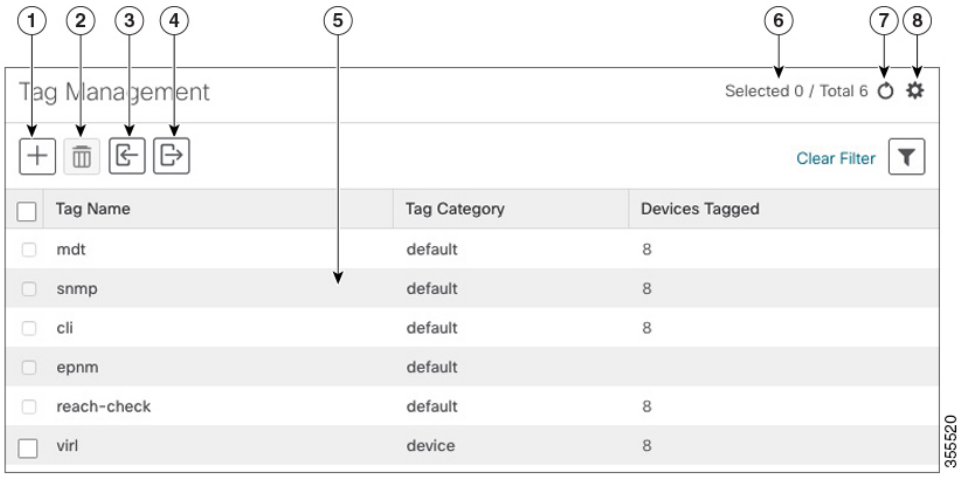

Item	Description
1	Click ⊞ to create new device tags. See Create Tags.
2	Click 🗑 to delete currently selected device tags. See Delete Tags.

Item	Description
3	Click to import the device tags defined in a CSV file into Cisco Crosswork Network Automation . See Import Tags, on page 61. You can also download a CSV file template by clicking this icon. The template includes sample data that you can use as a guide for building your own CSV file.
4	Click to export a CSV file that lists the tags that are currently configured and their attributes. You can update this file and import it back into Cisco Crosswork Optimization Engine to quickly add or edit multiple tags. See Export Tags, on page 63.
5	Displays the tags currently available in Cisco Crosswork Optimization Engine and their attributes.
6	Indicates the number of tags that are currently selected in the table.
7	Click to refresh the **Tag Management** window.
8	Click to choose the columns to make visible in the **Tag Management** window (see Set, Sort and Filter Table Data, on page 7).
	Click to set filter criteria on one or more columns in the **Tag Management** window.
	Click the **Clear Filter** link to clear any filter criteria you may have set.

Create Tags

You can create as many tags and tag categories as you want. If you will have many tags, it might be quicker to list them in a CSV file and import the file, instead of creating each tag individually. See Import Tags, on page 61.

Note Tag and tag category names are case-insensitive and can contain up to 128 alphanumeric characters, and can use full stops ("."), underscores ("_"), and hyphens ("-"). They cannot contain other special characters, symbols, or spaces.

Step 1 From the main menu, choose **Inventory Management** > **Tags**. The **Tag Management** window opens.

Step 2 Click . The **Create New Tags** pane opens.

Step 3 In the **Category** area:

- To associate your new tags with an existing category: Choose the category from the drop-down list.
- To associate your new tags with a new category: Click the **New Category** link, enter the new category's name in the text field, and click **Save**.

All the new tags you create after this step will be assigned to the category you selected or created.

Step 4 In the **Tags** area: Start entering the names of the new tags that you want to create. Press **Return** after you type each tag.

To keep from entering duplicate tags, click the **Show Tags** link. The **Create New Tags** window will list only the tags that already exist in your currently selected category.

Step 5 When you are finished entering new tags, click **Save**.

What to do next

Add tags to devices. See Apply or Remove Device Tags, on page 62.

Import Tags

Complete the steps below to create a CSV file that lists the tags you want to apply to your devices, and then import it into Cisco Crosswork Optimization Engine. This is the easiest way to create a lot of new tags and tag categories quickly.

When you import the CSV file, any tags not already in the database will be added. Tags with the same name as an imported tag will be overwritten. For this reason, it is a good idea to export a backup copy of all your current tags before import (see Export Tags, on page 63).

Step 1 From the main menu, choose **Inventory Management** > **Tags**.

Step 2 Click to open the **Import CSV File** dialog box.

Step 3 If you have not already created a CSV file to import:

a) Click the **Download sample 'Tags template (*.csv)' file** link and save the CSV file template to a local storage resource.

b) Open the template using your preferred tool. Begin adding rows to the file, one row for each tag. Use a comma to delimit each field within a row. Use a semicolon to separate multiple entries in the same field.

Field	Description	Required or Optional
Tag Name	Enter the name of the tag. For example: `SanFrancisco` or `Spine/Leaf`.	Required
Tag Category	Enter the tag category. For example: `City` or `Network Role`.	Required

Be sure to delete the sample data rows before saving the file, or they will be imported along with the data you want. The column header row can stay, as it is ignored during import.

c) When you are finished, save the new CSV file.

Step 4 Click **Browse** to navigate to the CSV file you just created and then click **Open** to select it.

Step 5 With the CSV file selected, click **Import**.

The tags and tag categories that you imported should now be displayed in the **Tag Management** window.

What to do next

Add tags to devices. See Apply or Remove Device Tags, on page 62.

Apply or Remove Device Tags

Tags and their categories are your main tool for grouping devices. Once you have tagged a set of devices with the same tag, they are considered part of a group, and you can manage them more easily.

In order to apply a tag to a device or group of devices, the tag must already exist (see Create Tags, on page 60).

You can apply a maximum of 15 tags to any one device.

To apply tags to a device or set of devices, do the following:

Step 1 From the main menu, choose **Inventory Management** > **Devices**. The **Network Devices** tab is displayed, showing the list of devices.

Step 2 (Optional) If the list is long, click to set one or more filters and narrow the list to only those devices you want to tag.

Step 3 Check the check box next to the device(s) you want to tag. If you select multiple devices, any changes you make will be applied to all the devices you selected.

Step 4 From the toolbar, click . The **Modify Tags** window opens, showing the tags currently applied to the device(s) you selected.

Step 5 Click in the **Type to autocomplete item** field to display the list of existing tags, or begin typing the name of the tag you want.

Step 6 Click on individual tags in the list to add them to the list of tags applied to the device(s). To delete an applied tag, click the X icon shown next to that tag.

Delete Tags

To delete device tags, do the following:

Note If the tag is mapped to any devices, then the tag cannot be deleted.

Step 1 Export a backup CSV file containing the tags you plan to delete (see Export Tags, on page 63).

Step 2 From the main menu, choose **Inventory Management** > **Tags**. The **Tag Management** window is displayed.

Step 3 Check the check box next to the tags you want to delete.

Step 4 From the toolbar, click .

Step 5 The confirmation dialog box will list the number of devices currently using the tag(s) you are about to delete. Click **Delete** to confirm deletion.

Export Tags

You can quickly export tags and tag categories to a CSV file. This will allow you to keep backup copies of your tags. You can also edit the CSV file as needed, and re-import it to overwrite existing tags. Note that you will need to re-associate devices and tags in some cases.

Step 1 From the main menu, choose **Admin** > **Tags**.

Step 2 (Optional) In the **Tag Management** window, filter the tag list as needed.

Step 3 Check the check boxes for the tags you want to export. Check the check box at the top of the column to select all the tags for export.

Step 4 Click . Depending on your browser, you will be prompted to select a path and file name to use when saving the CSV file, or to open it immediately.

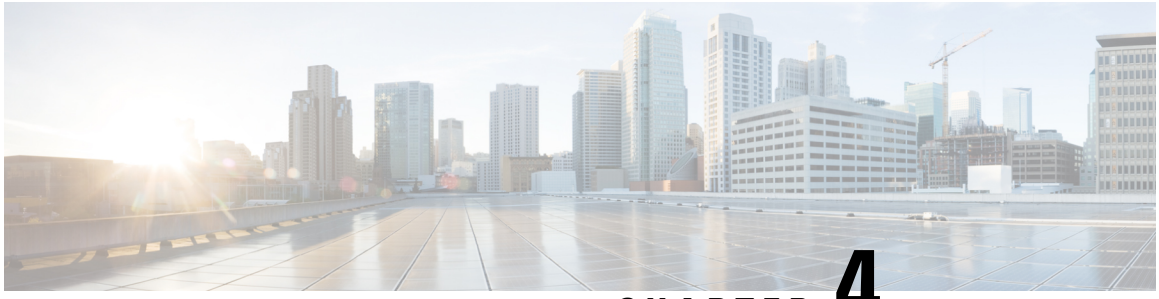

CHAPTER 4

Visualize the Network

Cisco Crosswork Optimization Engine provides a real-time graphical, topological map view of devices, links, and TE tunnels between them. This section focuses on device and link visualization features, and customizing the topology map.

For information on TE tunnel management and visualization, see Visualize SR Policies and RSVP-TE Tunnels, on page 83.

- Network Topology Map, on page 65
- Visualize Devices, on page 74
- Visualize Links, on page 78

Network Topology Map

The network topology can be displayed on a logical map or a geographical map, where the devices and links are shown in their geographic context. From the map, you can drill down to get detailed information about devices and links.

To get to the topology map, choose **Optimization Engine** from the left navigation bar, and click **Traffic Engineering**.

Figure 11: Network Topology Map - Devices and Links

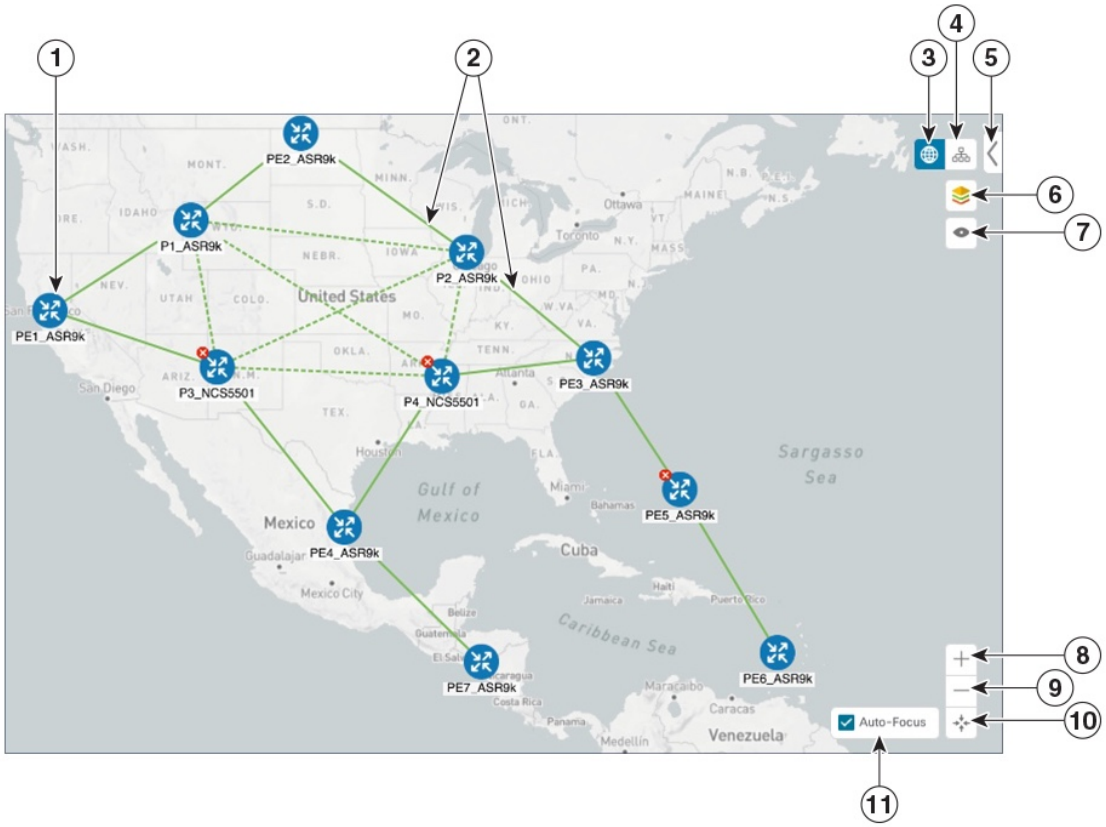

Callout No.	Description
1	**Single Device**: A blue device icon means the device is reachable; a gray icon means the device is not reachable. To view a device configuration summary, hover the mouse cursor over the device icon. A pop up window displaying the host name, state, node ID, and device type appears. To view device details, click on the device icon. The **Device Details** window appears to the right. See Get More Information About Devices on the Map, on page 74. **Device Cluster**: If devices are in close physical proximity, the geographical map shows them as a cluster. The number in a blue circle (2) indicates the number of devices in the cluster. Displaying devices in this manner helps prevent overlap and clutter on the map.
2	**Links**: A solid line indicates a *single link* between two devices. If there is more than one link between two devices, or between a device and a cluster of devices, the line is shown dashed instead. A dashed line indicates an *aggregated* link that represents more than one link.

Callout No.	Description
3	**Geographical Map**: Click this icon to view the geographical map. The geographical map shows single devices, device clusters, links, and tunnels, superimposed on a map of the world. Each device location on the map reflects the device's GPS coordinates (longitude and latitude) as defined in the device inventory.
4	**Logical Map**: Click this icon to toggle from the geographical map to the logical map. The logical map shows devices and their links, positioned according to an automatic layout algorithm, ignoring their geographical location. You can change the layout algorithm; see Change the Layout of a Logical Map, on page 69. The logical map displays up to 5000 devices and never displays devices in clusters. If you drill down to the logical map from a geographical cluster at the maximum zoom level, the logical map shows devices that are located in the same location. See Identify the Members of a Cluster, on page 78.
5	**Expand/Collapse/Hide Side Panel**: Expand or collapse the contents of the side panel. Close the side panel to get a larger view of the topology map.
6	**Display Preferences**: Lets you edit display settings for devices, links, and tunnel metrics. See Change Display Settings for Links, Devices, and TE Tunnel Metrics, on page 70.
7	**Custom Map View**: Lets you create a named custom view using the settings and layout for your current map, or display a custom view you have created previously. See Create Custom Map Views, on page 72. **Note** • The map zoom level will not be saved. • Your custom map views are not user-specific.
8	**Zoom In**: Click this icon to zoom in on the selected area; for example, to view clustered devices on the geographical map.
9	**Zoom Out:** Click this icon to zoom out from a selection area.
10	**Zoom Fit**: Lets you automatically scale the map to fit your zoom area.
11	**Auto-Focus**: When checked, automatically zooms in on selected tunnels. This option is selected by default. If you uncheck this option or navigate away from the map; it will revert to the default display.

Troubleshoot Network Topology Map

If you encounter topology issues, such as topology components not rendering as expected or component data not displaying on the map, Cisco recommends the following:

- If you cannot see geographical map tiles: Make sure your browser has Internet connectivity to your selected geographical map services vendor. The map services vendor and the verndor's URL are set by the system administrator, as explained in Configure Geographical Map Settings, on page 69.

- If your devices are missing from the geographical map: Ensure that latitude and longitude data was included when onboarding your devices, or entered later. Cisco Crosswork Optimization Engine cannot position devices properly on the geographical map without location information.
- Devices that do not have geographical coordinates default to 0° latitude and 0° longitude.
- If your devices are appearing in the wrong location on the geographical map, confirm that you have entered the latitude and longitude values in the correct order via the UI or in the CSV file you uploaded.
- If you are having intermittent problems displaying the map or your devices: Clear your browser cache and try again.

Device and Link Icons

The following tables describe the icons used to represent device states, link states, and device types in the Cisco Crosswork Optimization Engine user interface.

Table 5: Device State Icons

Icon	Description
	The device is reachable.
	The device is unreachable.
	The device has an unknown reachability state (its reachability cannot be determined).
	The device is operational.
	The device is not operational. It is either not up, or unreachable, or both. A number in a circle is shown next to this icon. The number indicates the number of recent errors and can be clicked on to display error details.
	Some connections to the device are down.

Table 6: Link State Icons

Icon	Description
	Link is down.
	Link is up and traffic is passing through it.
	Link is degraded. If some (but not all links) in an aggregated link are down, the aggregated link shows a degraded icon. The link will also show as degraded if only one direction of an L2 or L3 link was discovered instead of both directions. Click the degraded icon to see exactly which link or interface is down. If *all* links in an aggregated link are down, the connectivity link shows a link down icon.

Table 7: Device Icons

Icon	Description
Note For specific TE tunnel icons available on the topology map, see SR Policies and RSVP-TE Tunnels Topology Map, on page 89. If the device is unreachable, the icon is grey.	
	Router
	Device is reachable, but is undefined or of an unknown type

Configure Geographical Map Settings

The geographical map lets you position your network devices on a world map and monitor them within their geographical context. The displayed world map is imported by accessing the map vendor's site over the Internet (online mode). The look of the map will vary depending on the map vendor you choose.

By default, the client machine from where you access Optimization Engine UI is setup to get map tiles from a specific Mapbox URL over internet connection. If required, you can use a different map vendor (such as Google Maps or OpenStreetMap) by providing the appropriate URL. Both of these options require an Internet connection from your client machine.

Cisco Crosswork Optimization Engine administrator privileges are required to change these settings.

Step 1 From the main menu, choose **Admin** > **Visualization Settings**.

Step 2 Click the **Map** tab.

Step 3 From the **Map Provider** drop-down list, choose one of the following:

- **Mapbox**—Specifies that you want to display the geographical map using the default map provider.
- **Custom**—Identifies the map tiles source (using an Internet connection). To use a map provider other than Mapbox, you must provide the URL for map tiles access. Be sure to request the exact format of this URL from the map tiles provider.

Step 4 If you are using a custom map provider, in the **Map Source URL** field, enter the URL for map access.

Step 5 Click **Save**.

Step 6 Navigate to **Optimization Engine** > **Traffic Engineering** and confirm that the map is displayed correctly.

Change the Layout of a Logical Map

When you open the logical map, it is displayed according to the default standard layout. You can change the layout, but any changes you make will not persist if you close the map. To save your layout changes, create a custom view (see Create Custom Map Views, on page 72).

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering**.

Step 2 In the top-right corner of the map, toggle from the geographical map view to the logical map view.

Step 3 In the logical map, click the **System Layouts** icon in the toolbar to access the layout options.

Figure 12: System Layouts

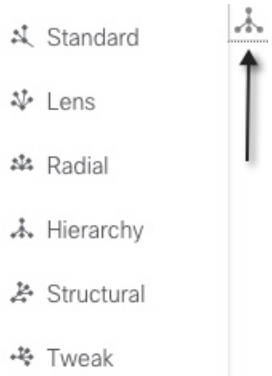

Step 4 Choose one of the predefined options to rearrange the devices and links in the map according to your preference:

- Standard (default)—Maintains consistent link length and distributes devices evenly. This ensures that adjacent devices are closer to each other and prevents overlap.
- Lens—Positions highly connected devices in the center, and moves less-connected devices out to the edges. This layout is especially useful in large networks.
- Radial—Arranges the devices in a circular style around the original subject. Each generation of devices becomes a new concentric ring that orbits the original parent. This layout is useful in networks where each parent has many child devices.
- Hierarchy—Displays devices in a family tree, where child devices are shown in horizontal layers underneath their parents.
- Structural—Groups devices with similar attributes together in a fan shape. This layout gives you an overview of the clusters in the network.
- Tweak—Adjusts the layout as the network evolves. As devices and links are added and removed, the layout adapts itself, allowing you to visualize network changes.

Change Display Settings for Links, Devices, and TE Tunnel Metrics

You can change the appearance of links and devices that are displayed on the topology map. You can also choose to show metrics on the topology map for selected TE tunnels. To do this click on the topology map.

Links

By default, colored link utilization thresholds and aggregated links are displayed. To change this select the **Links** tab and make the appropriate changes.

Figure 13: Links-Display Preferences

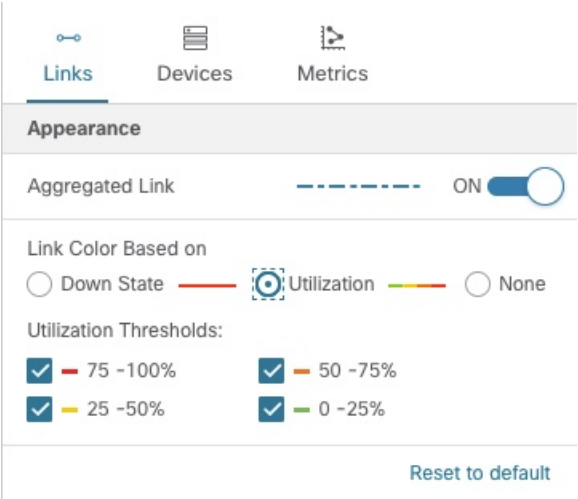

Devices

By default, the device host name and the state of the device is shown. To hide device states or display devices with a different label (Node IP, ISIS System ID, OSPF Router ID, or TE Router ID), select the **Devices** tab.

Figure 14: Devices-Display Preferences

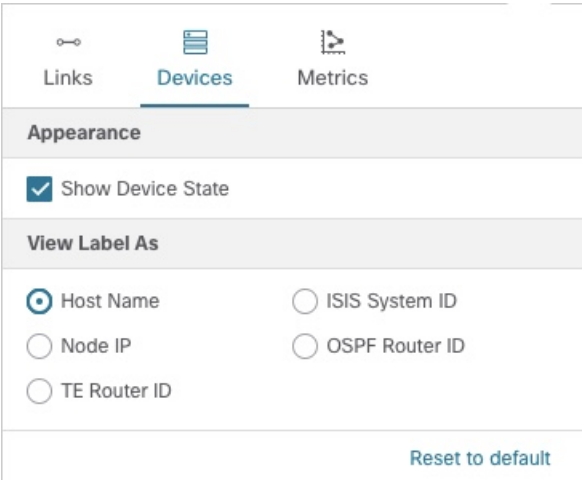

TE Tunnel Metrics

By default, TE tunnel metrics are not displayed. To display tunnel metrics, do the following:

1. Select tunnels from the **SR Policies** or **RSVP-TE Tunnels** table.
2. Click the **Show IGP Path** checkbox on the topology map.
3. Click and click the **Metrics** tab.
4. Check the metrics you want to see displayed.

Figure 15: TE Tunnel Metrics-Display Preferences

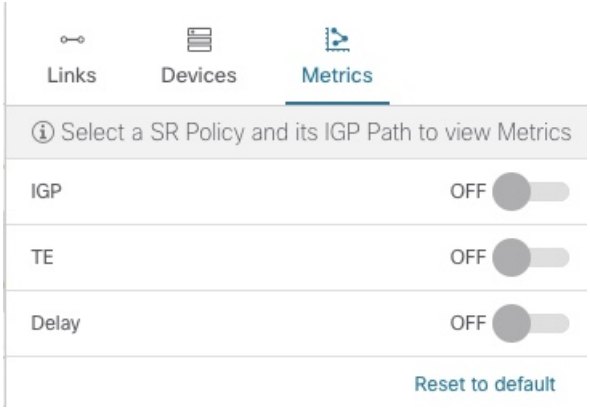

Create Custom Map Views

When you rearrange the devices and links on a map, your changes are not normally saved. When you open the map later, your map settings are lost.

To easily recreate a useful map layout, you can save it as a named custom view and quickly retrieve it, without having to rearrange the map each time. This is especially useful when managing large networks with many devices.

When you save a custom view, the following settings will be saved:

- Whether it is a geographical or logical map.
- Device positions in the logical map layout.
- Whether bandwidth utilization visualization is enabled or disabled.

Note

- The map zoom level will not be saved.
- Your custom map views are not user-specific. It is shared and can be modified by all users using the same Cisco Crosswork Optimization Engine server.

To create custom views:

Step 1 Choose **Optimization Engine** > **Traffic Engineering** from the left navigation bar.

Step 2 Customize the current map view until it contains only the information you want and until the layout meets your needs.

Step 3 When you have the view the way you want it, click .

Step 4 Click **Save View** and the Save View popup displays a new, blank input field under the **Name** field.

Step 5 Enter a unique name for the new custom view and click **Save**.

What to do next

Retrieve, update and delete your custom views as explained in Manage Custom Map Views, on page 73.

Manage Custom Map Views

You can display, update or delete any of the custom views created using the instructions in Create Custom Map Views, on page 72. This includes custom views created by other users.

To manage custom views:

Step 1 Open the topology map by choosing **Optimization Engine** > **Traffic Engineering** from the left navigation bar.

Step 2 Click and the following options are displayed:

- **View Saved Views**—Displays all saved custom views. You can choose to see only custom views saved with your user ID (My Views tab) or all custom views that have been saved on the server (All Views tab).
- **Save View**—Allows you to save the current view.
- **Save View As**—Click this option if you are currently modifying a custom view and want to save changes as a new custom view with a new name.
- **Rename View**—Click this option if you are currently modifying a custom view and want to rename it.

Step 3 To delete a custom view:

a) Click the **View Saved Views** to display the list of custom views.

b) Find the view you want to delete and click from the custom view.

Figure 16: Saved Custom Views

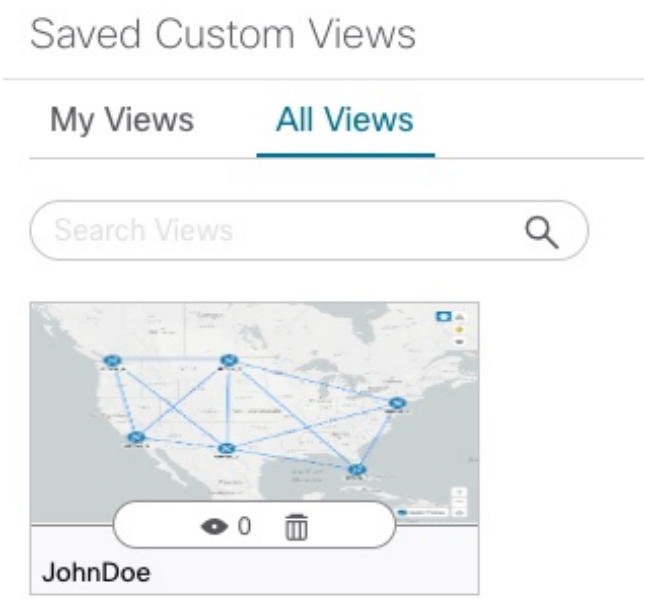

Visualize Devices

Cisco Crosswork Optimization Engine displays discovered devices in your network and gives you the ability to access the device via SSH or Telnet. To view link state icons, see Device and Link Icons, on page 68. This section contains the following topics:

- Get More Information About Devices on the Map, on page 74
- Access the Device Console, on page 77
- Identify the Members of a Cluster, on page 78

Get More Information About Devices on the Map

In the topology map, hover over a device icon to open a popup window with the most important device details: hostname, reachability state, IP address, and type. Click on the device icon to open the **Device Details** pop-up window, where you can view more detailed information about the device and its associated links. See the following examples.

Click on the device icon to open the **Device Details** pop-up window where you can view more detailed information about the device and its associated links. In a multiple IGP setup, you can view all the IGP, IS-IS, and OSPF processes. See the following examples:

Figure 17: Device Details Popups

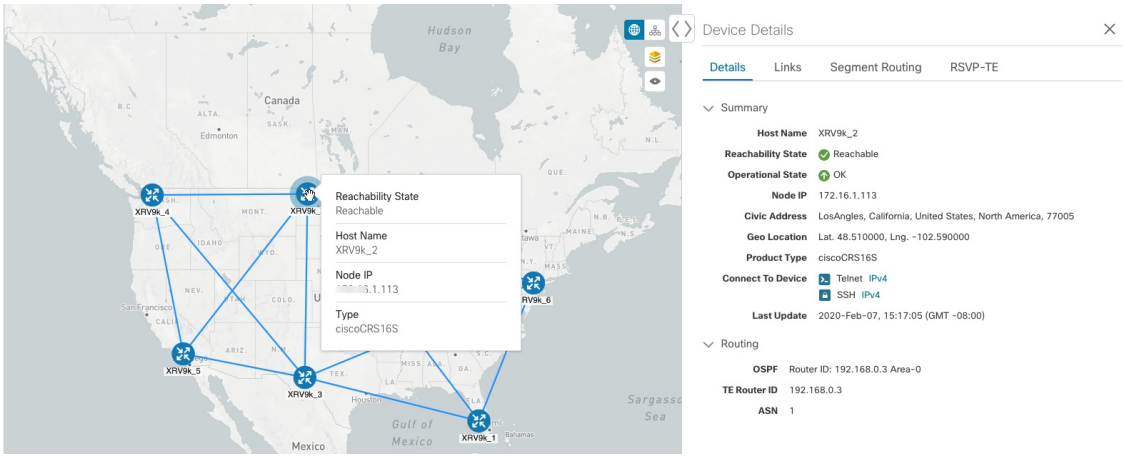

Figure 18: Multiple IGP: IGP Processes

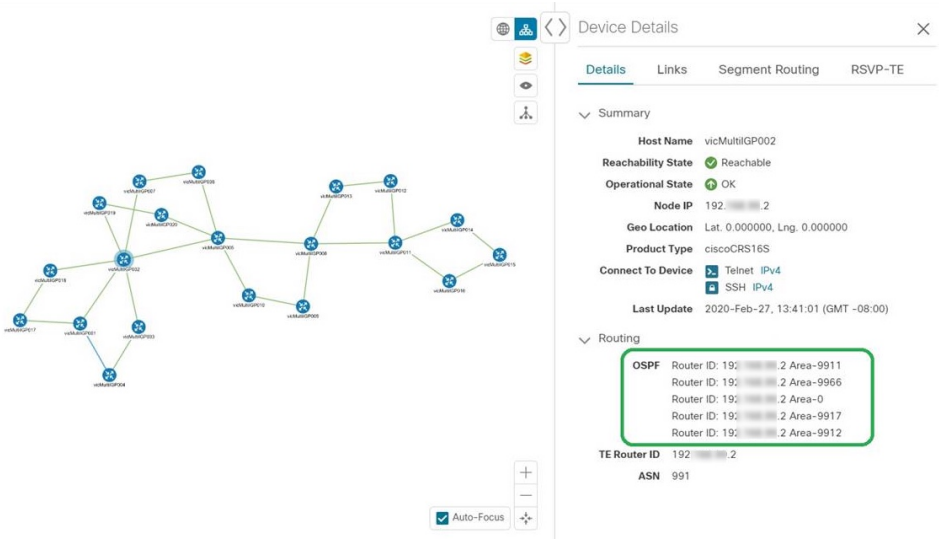

Figure 19: Multiple IGP: ISIS Processes

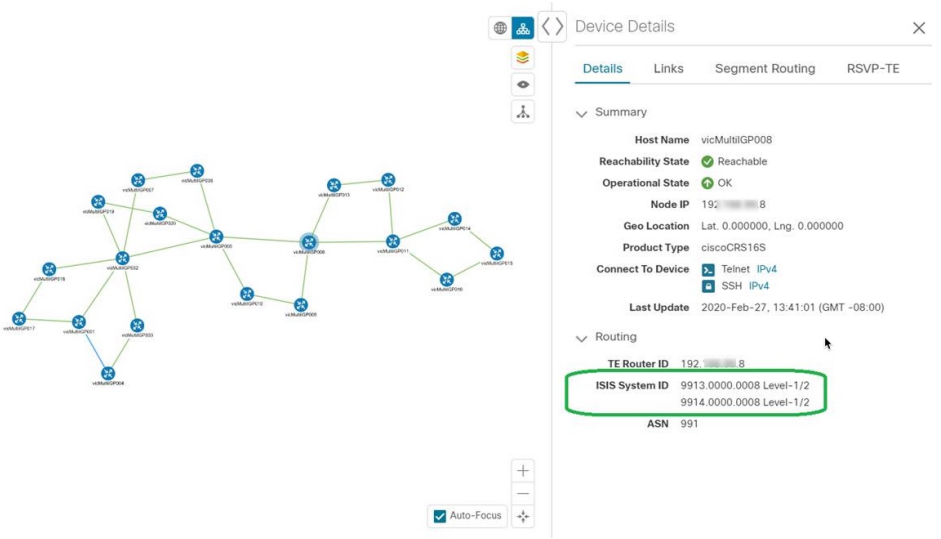

Figure 20: Multiple IGP: OSPF and ISIS Processes

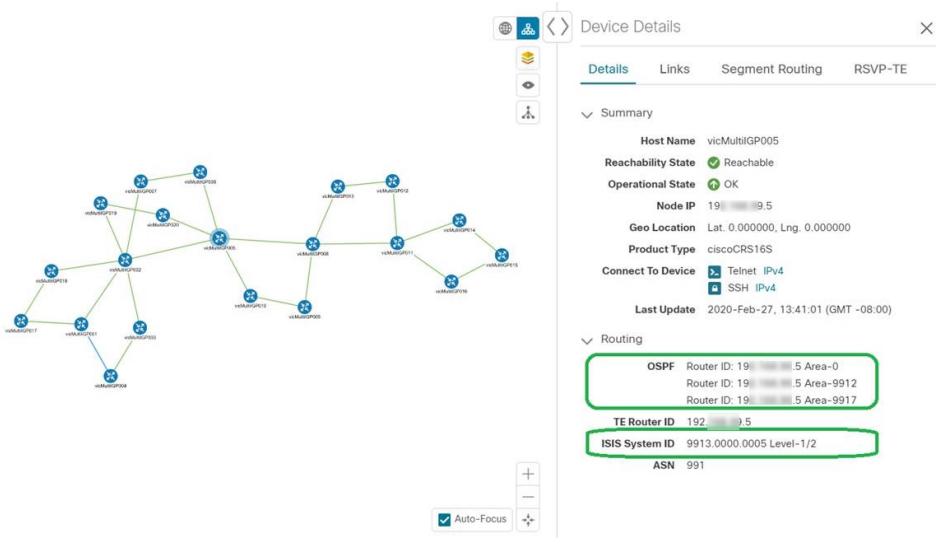

In the **Device Details** window, click on the **Links** tab to see a list of all of the device's links to other devices, as in the following example (see Get More Information About Links, on page 78):

Figure 21: Links Tab of Device Details Window

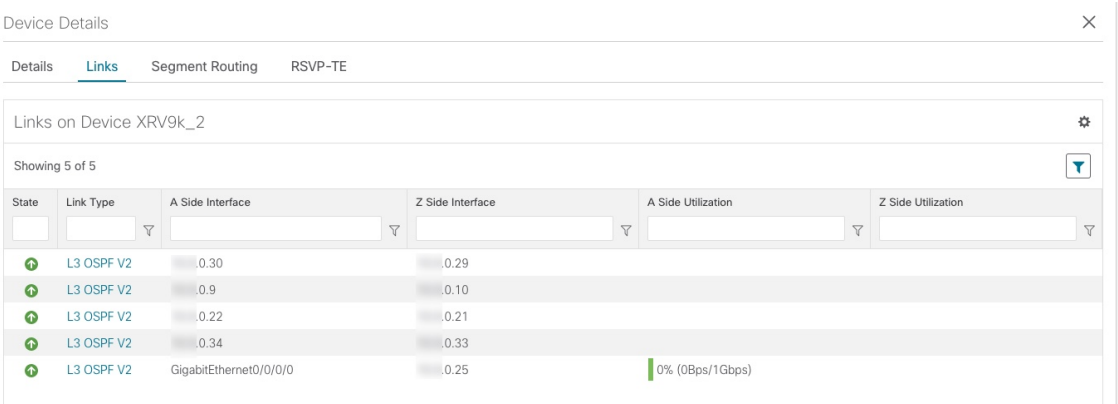

Access the Device Console

After drilling down to a device's details from the topology map, you can access the device's CLI command console from the **Device Details** window (see Get More Information About Devices on the Map, on page 74).

Before you begin

- Depending on your environment, your local machine may not have direct access to your network devices (for example: you cannot ping the device's management address directly from the command line on your local machine). If this is the case, you may need to configure a tunnel. Contact Cisco Services for assistance with this more advanced configuration.
- Be sure you have installed on your client an application that can connect to devices via Secure Shell (SSH) or Telnet.

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering**.

Step 2 In the topology map, click on the icon representing the device to which you want to connect. The **Device Details** window displays its **Details** tab, with the device hostname, reachability state, IP address, and other details.

Step 3 In the **Connect to Device** field, click the relevant link to connect to the device console via Telnet , or via SSH .

If you have already defined a default connectivity application that you want to launch, Cisco Crosswork Optimization Engine launches your selected application and attempts to connect to the device. Log into the device and enter the commands you want.

If you have not defined a default application to launch, your browser will prompt you to select one. Your choices and how they are presented will be appropriate for your client operating system, the applications you have installed, and the connectivity protocol you choose. Select the application you want and, for convenience, make sure that you select the check box indicating that this is your default choice before continuing.

Identify the Members of a Cluster

When there are multiple devices that are too close to be shown individually at the current Zoom level, they are combined together and shown as a single cluster. The cluster is represented on the geographical map by a circle with a number in its center, indicating the number of devices in the cluster.

Zoom in on a cluster to see the individual devices in the cluster displayed on the map.

If cluster members are very close to each other or in the same location, zooming in will not show the individual devices. In this case, follow these steps to see the individual members of the cluster:

Step 1 In the geographical map, click 2. The map zooms in on the cluster area.

Step 2 Click 2 again. If you are at the maximum zoom level, the geographical map toggles to the logical map and displays the individual devices in the cluster. When you close the view, you will be switched back to the geographical map.

Visualize Links

Cisco Crosswork Optimization Engine displays the links between devices and gives you the ability to configure and view bandwidth utilization on these links. To view link state icons, see Device and Link Icons, on page 68. This section contains the following topics:

- Get More Information About Links, on page 78
- Show Bandwidth Utilization for Links on the Map, on page 80
- Define Color Thresholds for Link Bandwidth Utilization, on page 81

Get More Information About Links

You can drill down in the topology map to view detailed information about links, using either of these methods:

- Click on an aggregated link (symbolized by a dashed line) to show the individual links in the side panel.
- Click on a single link (solid line) to show the **Link Details** page.

The **Links** page provides information about the configuration and status of all of a device's links, including each link's type, interfaces, and utilization (you can get the same information from the **Links** tab on the **Device Details** window; see Get More Information About Devices on the Map, on page 74). The **Links** window lists all the underlying links in the aggregation, as in the following example:

Figure 22: Links Window

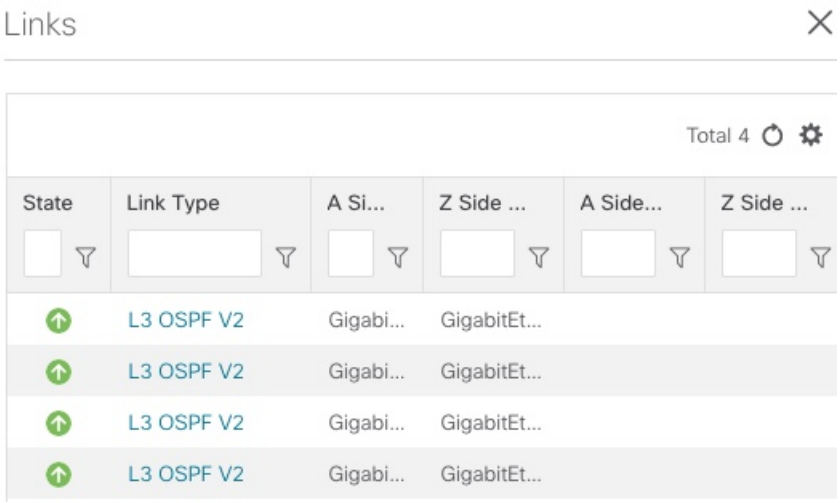

Use the expand and collapse icons (< and >) to the left of the **Links** title to expand the window to the entire screen, or collapse the window back to its normal size.

Click ⚙ to choose the columns to make visible in the **Links** window's table:

- State—Displays each link's state: up, down, degraded, and so on (see Device and Link Icons, on page 68).
- Link Type—Displays the type of link. Click on the link type to open the **Link Details** window for the specific link.
- A Side Device—Displays the originating device for the link.
- Z Side Device—Displays the destination device for the link.
- A Side Interface—Displays the originating interface for the link.
- Z Side Interface—Displays the destination interface for the link.
- A Side Utilization—Displays the percentage of bandwidth consumption on the originating side of the link.
- Z Side Utilization—Displays the percentage of bandwidth consumption on the destination side of the link.

You can also use sorts and filters in the **Links** window to focus the table on only the links in which you are interested (see Set, Sort and Filter Table Data, on page 7).

The **Link Details** window provides information about the configuration and status of a single link, including link type, the link's interfaces, associated adjacent segment IDs, and so on. You can also find information about any associated TE tunnels by clicking on the Segment Routing and RSVP-TE tabs.

Figure 23: Link Details

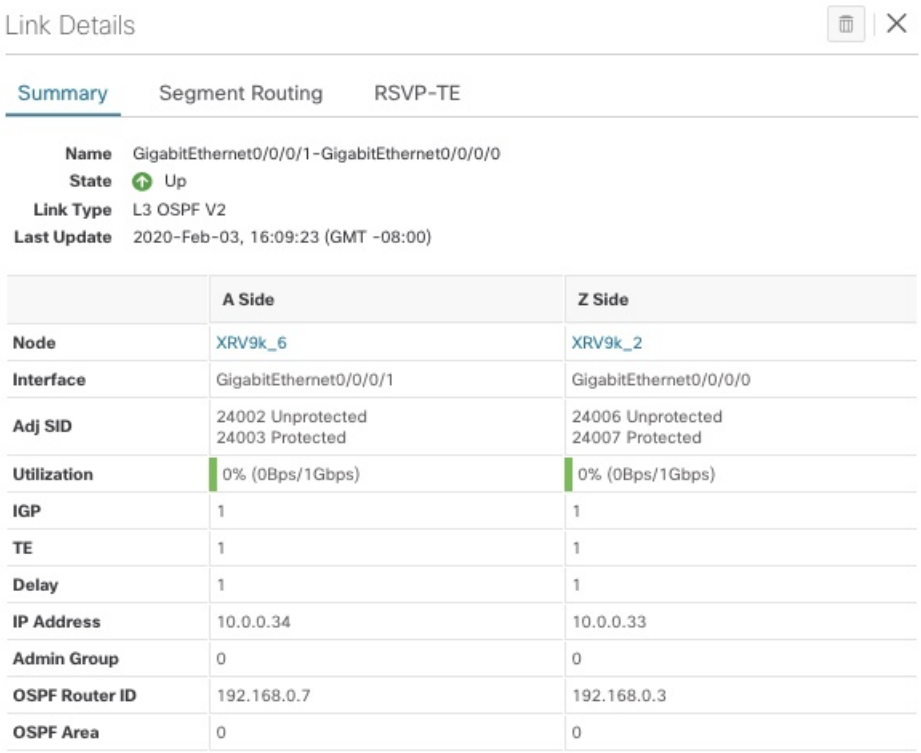

	A Side	Z Side
Node	XRV9k_6	XRV9k_2
Interface	GigabitEthernet0/0/0/1	GigabitEthernet0/0/0/0
Adj SID	24002 Unprotected 24003 Protected	24006 Unprotected 24007 Protected
Utilization	0% (0Bps/1Gbps)	0% (0Bps/1Gbps)
IGP	1	1
TE	1	1
Delay	1	1
IP Address	10.0.0.34	10.0.0.33
Admin Group	0	0
OSPF Router ID	192.168.0.7	192.168.0.3
OSPF Area	0	0

Show Bandwidth Utilization for Links on the Map

In the geographical map and in the logical map, you can enable visualization of the bandwidth utilization for links over which circuits are provisioned. When bandwidth utilization visualization is enabled, links in the map are colored based on the percentage of total bandwidth currently utilized on the link. The utilization value is a percentage calculated by dividing link traffic by link capacity.

In this way, you can easily identify when a link is over-utilized or approaching over-utilization. Bandwidth visualization is enabled by default. The color of the link indicates the percentage of total bandwidth being used by provisioned circuits on the link:

- Green—0–25% usage
- Yellow—25–50% usage
- Orange—50–75% usage
- Red—75–100% usage

You can adjust the thresholds for each color as needed (see Define Color Thresholds for Link Bandwidth Utilization, on page 81). When visualization is disabled, the links are shown only in blue.

Please note that link bandwidth utilization data can be collected and displayed only if the linked devices are added to and managed in the device inventory.

To enable or disable visualization of bandwidth utilization:

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering**.

Step 2 In the top-right corner of the map, click , to open display preferences which can be used to toggle the display of bandwidth utilization. When usage visualization is enabled, the links can be shown in green, yellow, orange, or red, depending on their utilization. If you see only blue links, usage visualization is disabled. This option is selected by default. If you uncheck this option, navigate away from the map, and later return to the map; it will revert to the default option.

Define Color Thresholds for Link Bandwidth Utilization

Cisco Crosswork Optimization Engine comes with a default set of bandwidth utilization thresholds (percentage ranges) and corresponding color indicators. You can customize these to meet your needs, taking into account the following notes and limitations:

- You can enter values in the "To" ranges. Each row begins automatically from the end of the previous row's range.
- The thresholds must be sequential, meaning that each row's range must follow on from the previous row's range. For example, if the range in the first row is 0-25%, the second row's range must end with a value greater than 25.
- You cannot use the same color for multiple thresholds. For example, you cannot choose **Green** for both the first and second rows.

Administrator privileges are required to change these settings.

Step 1 From the main menu, choose **Admin** > **Visualization Settings**.

Step 2 Click the **Bandwidth Utilization** tab.

Step 3 In the **Polling Interval** field, enter a whole number from 5 to 60 (minutes) to specify how often links will be polled for bandwidth utilization. By default, link bandwidth is polled every 5 minutes.

Step 4 In the **Link Coloring Thresholds** area, define the criteria for coloring the links. Each row defines a color and the bandwidth percentage range that the color will represent. The default thresholds are:

- Green—0–25% usage
- Yellow—25–50% usage
- Orange—50–75% usage
- Red—75–100% usage

Step 5 Click **Save**.

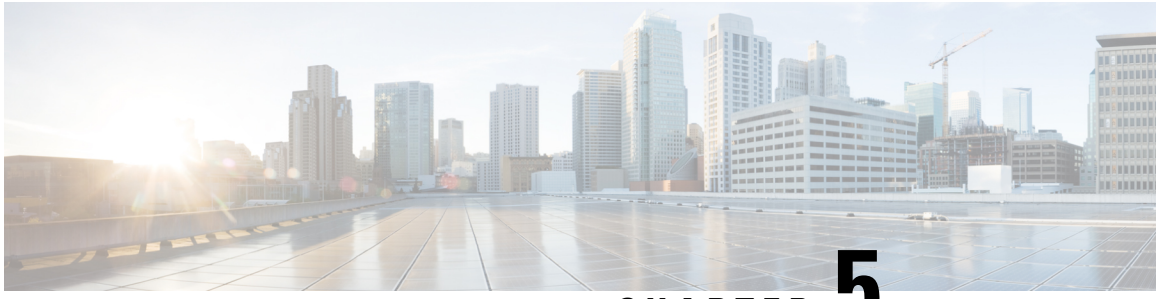

CHAPTER 5

Visualize SR Policies and RSVP-TE Tunnels

Note Throughout this document TE tunnels refer to both SR policies and RSVP-TE tunnels.

Cisco Crosswork Optimization Engine visualization provides the most value by giving you the ability to easily view and manage SR policies and RSVP-TE tunnels. By visually examining your network, the complexity of provisioning and managing these TE tunnels is significantly reduced.

To view supported TE tunnel features and limitations, see SR Policy and RSVP-TE Tunnel Support, on page 83.

This section contains the following topics:

- SR Policy and RSVP-TE Tunnel Support, on page 83
- SR Policy and RSVP-TE Tunnel Configuration Sources, on page 88
- SR Policies and RSVP-TE Tunnels Topology Map, on page 89
- Highlight a TE Tunnel on the Map, on page 92
- Show Participating Nodes and Links, on page 93
- Show IGP, Delay, and Traffic Engineering Metrics, on page 93
- SR Policies Table, on page 94
- RSVP-TE Tunnels Table, on page 96
- Visualize SR Policies and RSVP-TE Tunnels, on page 98
- Configure Affinity Mapping, on page 103
- Preview Disjoint SR Policies and RSVP-TE Tunnels, on page 104
- View TE Tunnels Belonging to a Disjoint Group, on page 107
- Create and Manage SR Policies, on page 107
- Create and Manage RSVP-TE Tunnels, on page 117

SR Policy and RSVP-TE Tunnel Support

The following lists provide an overview of SR policy and RSVP-TE tunnel supported and unsupported features. Contact your Cisco Crosswork Optimization Engine representative for any capabilities that are not documented in the following lists.

Table 8: Supported Features

Category	Capability	Notes
RSVP-TE	• PCE-initiated tunnels (provisioned or discovered by Crosswork Optimization Engine) • PCC-initiated tunnels (discovered by Crosswork Optimization Engine)	—
	• ERO strict hops • ERO loose hops (PCC-initiated only)	—
	FRR protection on tunnels provisioned by Crosswork Optimization Engine	—
	Path optimization objective min-metric (IGP,TE, or Latency)	—
	Path constraints (affinity and disjointness)	Only 2 RSVP tunnels per disjoint group or sub-id is supported
	Binding Label for explicit and dynamic tunnels	—
	Signaled Bandwidth	—
	Setup/Hold Priority	—

Category	Capability	Notes
SR Policy	• PCE-initiated tunnels (provisioned or discovered by Crosswork Optimization Engine) • PCC-initiated tunnels (discovered by Crosswork Optimization Engine) • SR On-Demand Next Hop (ODN) policies discovered by Crosswork Optimization Engine	—
	Single consistent Segment Routing Global Block (SRGB) configured on routers throughout domain covered by Crosswork Optimization Engine	If index SIDs are used and there are different SRGB bases along a path of a policy, the label can change along the path.
	• Prefix SID • Adjacency SID • EPE adjacency SID	—
	Protected and Unprotected adjacency SIDs	—
	Regular and Strict prefix SIDs	—
	SR policy optimization objective min-metric (IGP, TE, and Latency)	—
	SR policy path constraints (affinity and disjointness)	Only 2 SR policies per disjoint group or sub-id are supported
	Binding SID for explicit or dynamic policies	—
	Profile ID	—

Table 9: Unsupported Features and Limitations

Category	Description	Notes
RSVP-TE	Configuring loose hop ERO in COE	Only strict hops can be configured. If strict hops are not configured for every hop along the path and those hops are not remote interface IPs or loopback IPs, unexpected behavior may occur. For example, a tunnel may remain operationally down, hops may be modified, and so on.
	Named tunnels configured on PCCs	These tunnels are not discovered byCrosswork Optimization Engine.
	Tunnels with Loopback IPs other than TE router ID for headend or endpoint and path hops	—
	Display of active FRR protected paths in the topology map.	Crosswork Optimization Engine discovers FRR tunnels which are displayed in the topology map, but will not associate an actively protected tunnel with the FRR tunnel being used. The path in the topology map will not include FRR protected paths when protection is active.
	P2MP tunnels	—

Category	Description	Notes
SR Policy	Provisioning multiple candidate paths via Crosswork Optimization Engine	These paths are not discovered if configured on PCC. Crosswork Optimization Engine does not support configuration of these paths.
	Weighted Equal-Cost Multipath (WECMP)	—
	Multiple segment lists per candidate path	• Crosswork Optimization Engine does not support this configuration • If configured on a PCC, Crosswork Optimization Engine will not discover these segment lists.
	Visualization of multiple candidate paths	Only the current active path can be seen in the UI.
	Binding SIDs as Segment List Hops	—
	SR IGP Flexible Algorithm (Flex Algo)	—
	Anycast SIDs	—
	Hop count metric type for policies	Crosswork Optimization Engine does not support provisioning with this metric type and does not discover this metric type if configured on the PCC
	Routers that are not SR-capable	The assumption is that all routers discovered by Crosswork Optimization Engine are SR-capable
	SR policies with Loopback IPs other than TE router ID for headend/endpoint and prefix SIDs in segment list	—
	SR policy provisioned with IPv6 endpoints/hops	—
	SRv6	Only 2 SR policies per disjoint group/sub-id
	SR policy optimization objective min-metric with margin	Not supported for policies provisioned by Crosswork Optimization Engine. Margin is not discovered for PCC-initiated policies.

Category	Description	Notes
	SR policy constraints (resource exclusion or metric bound)	Not supported for policies provisioned by Crosswork Optimization Engine. Constraints are not discovered for PCC-initiated policies.

SR Policy and RSVP-TE Tunnel Configuration Sources

SR policies and RSVP-TE tunnels discovered and reported by Cisco Crosswork Optimization Engine may have been configured from the following sources:

- For SR policies:
  - PCC initiated—Policies configured on a PCC (see PCC-Initiated SR Policy Example, on page 88).
  - PCE initiated—Policies configured on PCE or policies created dynamically by Cisco Crosswork Optimization Engine. A TE tunnel that is configured using Cisco Crosswork Optimization Engine is the only type of TE tunnel that Cisco Crosswork Optimization Engine can modify or delete (see Create and Manage SR Policies, on page 107 or Create and Manage RSVP-TE Tunnels, on page 117).
- For RSVP-TE tunnels:

  Note These tunnels cannot be configured directly on a PCE.

  - PCC initiated—Policies configured on a PCC (see PCC-Initiated RSVP-TE Tunnel Example, on page 89 and Path Computation Client (PCC) Support, on page 41).
  - Dynamically created.

PCC-Initiated SR Policy Example

The following example shows a configuration of an SR policy at the headend router. The policy has a dynamic path with affinity constraints computed by the headend router. See SR configuration documentation for your specific device to view descriptions and supported configuration commands (for example: *Segment Routing Configuration Guide for Cisco ASR 9000 Series Routers*).

```
segment-routing
 traffic-eng
  policy foo
   color 100 end-point ipv4 1.1.1.2
   candidate-paths
    preference 100
     dynamic
      metric
       type te
      !
     !
     constraints
```

```
      affinity
       exclude-any
        name RED
       !
      !
     !
    !
   !
```

PCC-Initiated RSVP-TE Tunnel Example

The following is a sample device configuration for a PCC-initiated RSVP-TE tunnel. See the appropriate documentation to view descriptions and supported RSVP-TE tunnel configuration commands for your particular device (for example: *MPLS Command Reference for Cisco NCS 5500 Series, Cisco NCS 540 Series, and Cisco NCS 560 Series Routers*).

```
interface tunnel-te0
 ipv4 unnumbered Loopback0
 destination 172.16.255.5
 path-option 10 dynamic
!
```

SR Policies and RSVP-TE Tunnels Topology Map

To get to the topology map, choose **Optimization Engine** from the left navigation bar, and click **Traffic Engineering**.

For information on topology issues, or using the map to get information about devices and links, see Network Topology Map, on page 65 and Troubleshoot Network Topology Map, on page 67.

The following example shows the topology map with SR policies highlighted.

Figure 24: SR Policies Topology Map Example

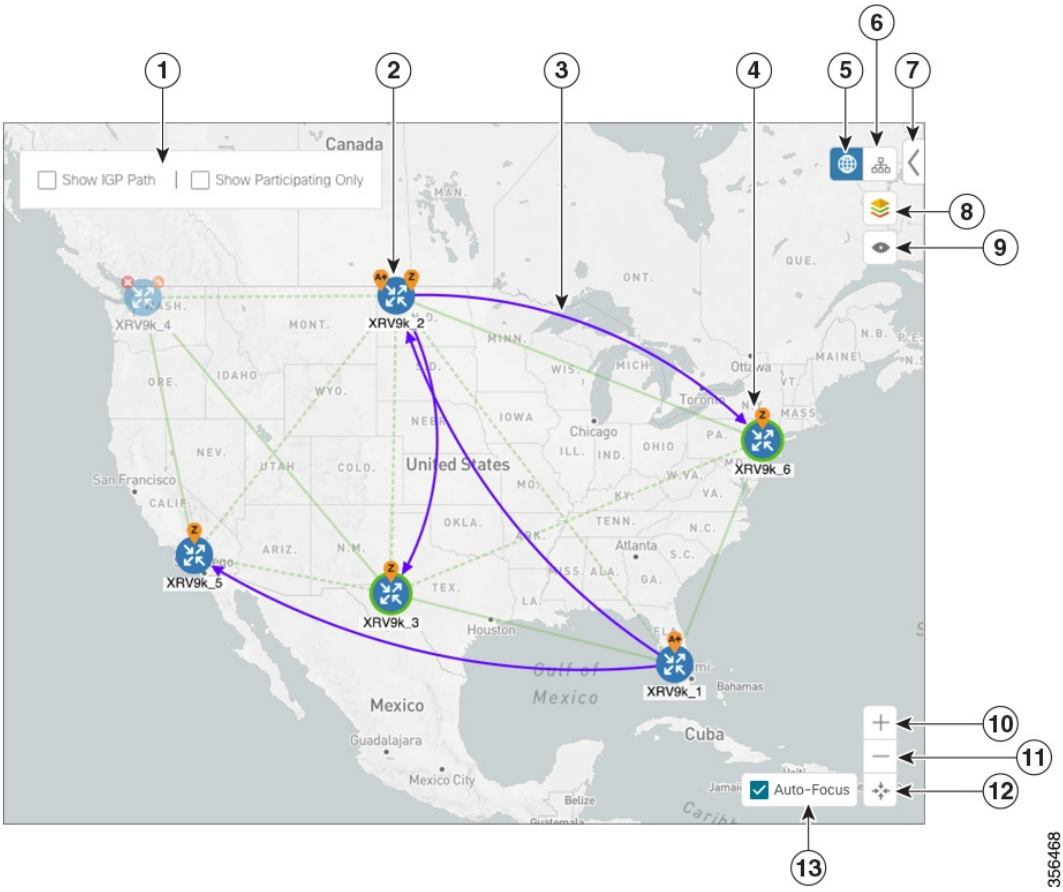

The display of RSVP TE tunnels is similar except for the following:

- The **Show IGP Path** option is not available.
- Record Route Object (RRO) paths are shown as straight lines.
- Explicit Route Object (ERO) paths are shown as curved lines.

Note If both RRO and ERO paths are available, the RRO path is displayed by default.

Figure 25: RSVP-TE Tunnels

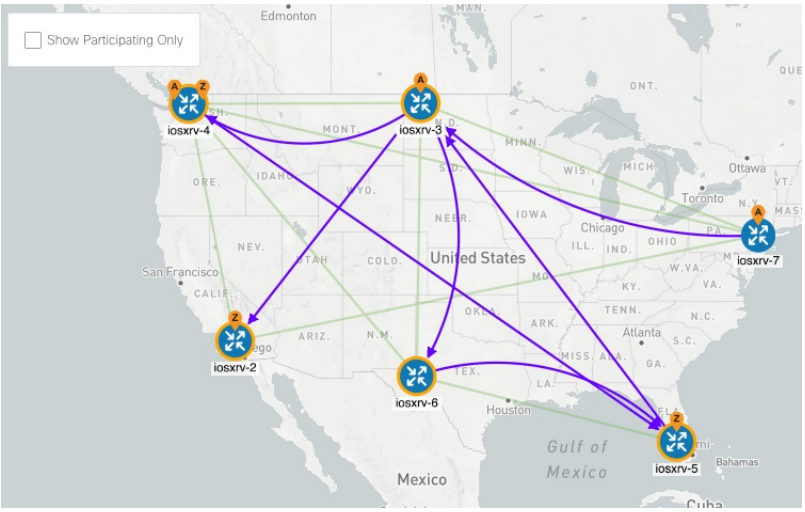

Callout No.	Description
1	Click the appropriate check box to enable the following options: • **Show IGP Path**—Displays the IGP path for the selected SR policy. This option is not available when viewing RSVP TE tunnels. • **Show Participating Only**—Displays only links that belong to selected TE tunnels. All other links and devices disappear.
2	**SR Policy and RSVP-TE Tunnel Origin and Destination**: If both **A** and **Z** are displayed in a device cluster, at least one node in the cluster is a source and another is a destination. The **A+** denotes that there is more than one SR policy or RSVP-TE tunnel that originates from a node. The **Z+** denotes that the node is a destination for more than one TE tunnel.
3	**SR Policies and RSVP-TE Tunnels**: When SR policies or RSVP-TE tunnels are selected from the SR Policies Table, on page 94 or RSVP-TE Tunnels Table, on page 96, they show as purple directional lines on the map indicating source and destination. An adjacency segment ID (SID) is shown as a green dot on a link along the path ( ).
4	SR Policies—A device with a green ( ) outline indicates there is a node SID associated with that device or a device in the cluster. RSVP-TE Tunnels—A device with a solid orange outline ( )indicates that it is a strict hop. A dashed orange outline indicates that a loose hop was discovered. **Note** RSVP-TE tunnels cannot be configured with loose hops when using the Crosswork Optimization Engine UI.

Callout No.	Description
5	**Geographical Map**: Click this icon to view the geographical map. The geographical map shows single devices, device clusters, links, and TE tunnels, superimposed on a map of the world. Each device location on the map reflects the device's GPS coordinates (longitude and latitude) as defined in the device inventory.
6	**Logical Map**: Click this icon to toggle from the geographical map to the logical map. The logical map shows devices and their links, positioned according to an automatic layout algorithm, ignoring their geographical location. You can change the layout algorithm; see Change the Layout of a Logical Map, on page 69. The logical map displays up to 5000 devices and never displays devices in clusters. If you drill down to the logical map from a geographical cluster at the maximum zoom level, the logical map shows devices that are located in the same location. See Identify the Members of a Cluster, on page 78.
7	**Expand/Collapse/Hide Side Panel**: Expand or collapse the side panel to see the full and truncated versions of the right-side panel. Close the side panel to get a larger view of the topology map.
8	**Display Preferences**: Lets you edit display settings for devices, links, and SR policy and RSVP-TE tunnel (with RRO) metrics. See Change Display Settings for Links, Devices, and TE Tunnel Metrics, on page 70.
9	**Custom Map View**: Lets you create a named custom view using the settings and layout for your current map, or display a custom view you have created previously. See Create Custom Map Views, on page 72.
10	**Zoom In**: Click this icon to zoom in on the selected area; for example, to view clustered devices on the geographical map.
11	**Zoom Out:** Click this icon to zoom out from a selection area.
12	**Zoom Fit**: Lets you automatically scale the map to fit your zoom area.
13	**Auto-Focus**: Zooms in on selected TE tunnels. This option is selected by default. If you uncheck this option, navigate away from the map, and later return to the map; it will revert to the default option.

Highlight a TE Tunnel on the Map

When many SR policies or RSVP-TE tunnels are displayed on the map, it may be difficult to view a particular path. To highlight a particular TE tunnel path on the map, navigate to **Optimization Engine** > **Traffic Engineering** > **SR-TE / RSVP-TE** tab, and hover over the SR policy or RSVP-TE tunnel. Prefix SID information will display under the node if it is part of the highlighted path.

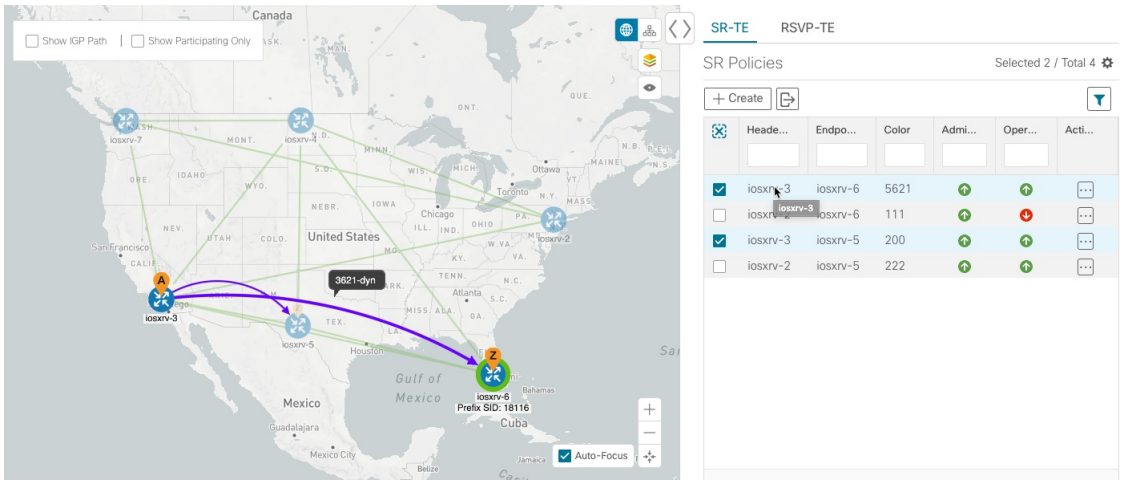

Show Participating Nodes and Links

To view only the nodes and links that are part of selected TE tunnels, do the following:

Step 1 From the **SR Policies** or **RSVP-TE Tunnels** table, select the TE tunnel you are interested in.

Step 2 From the top left box in the topology map, check the **Show Participating Only** check box.

Show IGP, Delay, and Traffic Engineering Metrics

Each link is assigned a metric value. The distance between two nodes is the sum of all the metric values of links along a path. To view IGP, Delay, or Traffic Engineering (TE) metrics on the topology map:

Step 1 Navigate to **Optimization Engine** > **Traffic Engineering**.

Step 2 Click the **SR-TE** or **RSVP-TE** tab and check the checkboxes next to the TE tunnels you are interested in. The TE tunnels are highlighted in the topology map.

Step 3 If viewing SR policies from the topology map, check the **Show IGP Path** checkbox.

Step 4 Click .

Step 5 Click the **Metrics** tab.

Step 6 Check the applicable metric check boxes you want displayed.

What to do next

To configure a TE tunnel based on one of these metrics, see Create Dynamic Path SR Policies, on page 110 or Create Dynamic Path RSVP-TE Tunnels, on page 118.

SR Policies Table

To get to the **SR Policies** table, choose **Optimization Engine** from the left navigation bar, and click **Traffic Engineering**. You will see the topology map and, to the right of the map, click the **SR-TE** tab.

Figure 26: SR Policies Table

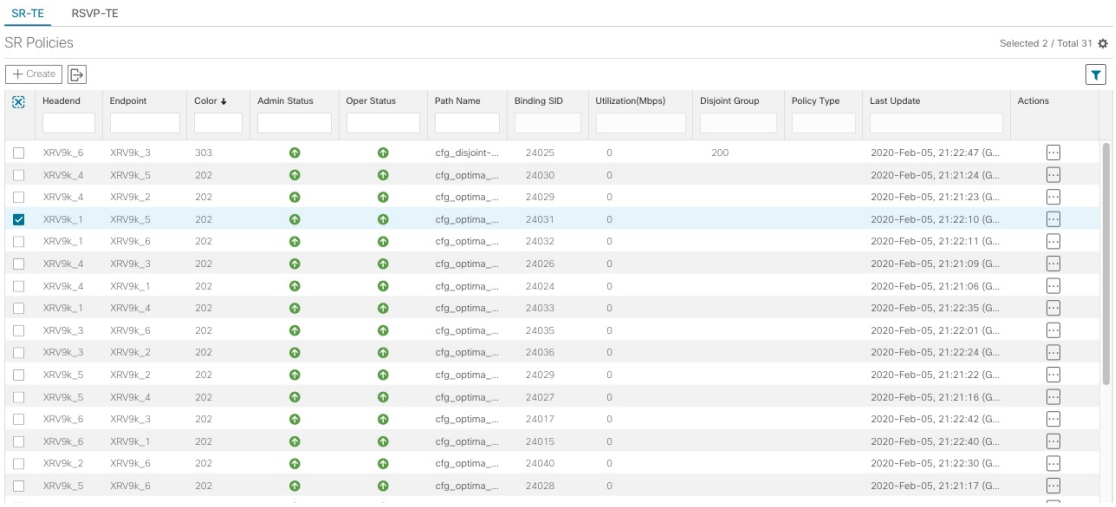

The **SR Policies** table provides the following functions:

- Display a list of all SR Policies discovered from the network.
- Configure new SR policies.
- Export the list of SR policies (click).
- Highlight SR policies on the map when selected from the table. To clear all selected policies, click .
- View SR policy details (click). See Get More Information About an SR Policy, on page 114). From the SR Policy Details page you can edit SR policies. However, only SR policies created from Crosswork Optimization Engine can be modified or deleted on the Crosswork Optimization Engine UI.
- Refresh ( ) the table or policy details (if in the **SR Policy Details** table). You can also view the date and time as to when the last refresh occured.

Note When creating or modifying SR policies, the refresh and auto-refresh functions are disabled in the tables.

The following information is available in the **SR Policies** table:

Note

- If a hostname is not available, click ⚙ and check the Headend IP and Endpoint IP checkboxes to show the respective IP addresses.
- Some fields may be blank depending on the SR policy type.

Table 10:

Column Heading	Description
Headend	Where the SR policy is instantiated.
Endpoint	The destination of the SR policy.
Color	A numerical value that distinguishes between two or more policies to the same node pairs (Headend – Endpoint). Every SR policy between a given headed and endpoint must have a unique color.
Admin Status	Administrative status of the SR policy.This is the status defined by the user.
Oper Status	Operational status of the SR policy. This is the state of the policy as reported by the system. For example, the user can define the Admin status as Up. However, if the policy is operationally down due to some network issues, then the Oper Status will display as Down.
Path Name	Name of SR policy path.
Binding SID	The binding segment is a local segment identifying an SR policy. Each SR policy is associated with a binding segment ID (BSID).
Utilization	Percentage of total bandwidth being used.
Disjoint Group	If applicable, the disjoint group the SR policy belongs in.
Policy Type	• Bandwidth Optimization • Bandwidth on Demand • Explicit • Dynamic
Last Update	Time when the most recent update for the policy was received from the network.

Column Heading	Description
Actions	Click ⋯ to Get More Information About an SR Policy, on page 114.

RSVP-TE Tunnels Table

To get to the **RSVP-TE Tunnels** table, choose **Optimization Engine** from the left navigation bar, and click **Traffic Engineering**. You will see the topology map and, to the right of the map,select the **RSVP-TE** tab.

Figure 27: RSVP-TE Tunnels Table

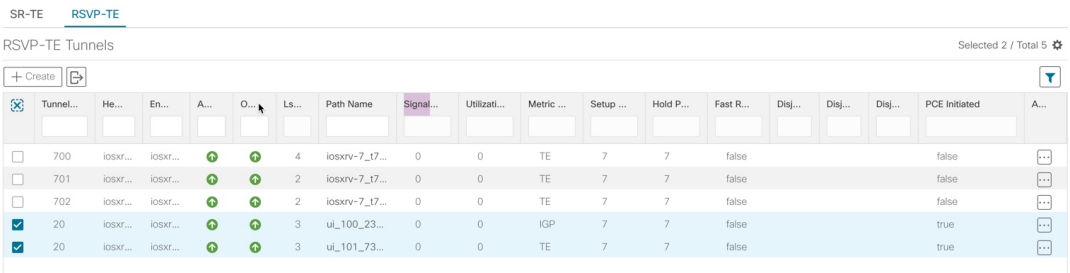

The **RSVP-TE Tunnels** table provides the following functions:

- Displays a list of all RSVP-TE tunnels discovered from the network.
- Configure new RSVP-TE tunnels.
- Edit RSVP-TE tunnels created using Crosswork Optimization Engine (click ⋯).

  Note Only tunnels created from Crosswork Optimization Engine can be modified or deleted on the Crosswork Optimization Engine UI.

- Highlight RSVP-TE tunnels on the map when selected from the table. To clear all selected tunnels, click ⊠.
- View RSVP-TE tunnel details (click on ⋯ link). See Get More Information About an RSVP-TE Tunnel, on page 121.
- Refresh (↻) the table. You can also view the date and time as to when the last refresh occurred.

  Note When creating or modifying RSVP-TE tunnels, the refresh and auto-refresh functions are disabled in the tables.

The following information is available in the **RSVP-TE Tunnels** table:

Note If a hostname is not available, click ⚙ and check the Headend IP and Endpoint IP checkboxes to show the respective IP addresses.

Table 11: RSVP-TE Tunnels

Column Heading	Description
Tunnel ID	The assigned tunnel ID value. The tunnel ID range is taken from the headend configuration.
Headend	Where the RSVP-TE tunnel is instantiated.
Endpoint	The destination of the RSVP-TE tunnel.
Admin Status	Administrative status of the RSVP-TE tunnel. This is the status defined by the user.
Oper Status	Operational status of the RSVP-TE tunnel. This is the state of the policy as reported by the system. For example, the user can define the Admin status as Up. However, if the policy is operationally down due to some network issues, then the Oper Status will display as Down.
LSP ID	This value is updated when there is a change to the path.
Utilization	Percentage of total bandwidth being used.
Metric Type	Type of metric (IGP, TE, or Delay).
Setup Priority	There are 8 (0 - 7) setup priorities. 0 is the most preferred. The setup priority is used to define preference for preempting less preferred tunnels. The most preferred tunnels can push the other less preferred tunnels out of the way.
Hold Priority	There are 8 (0 - 7) hold priorities. The holding priority is used to define a priority maintaining the currently established tunnel. You can have a tunnel that you never want go down, but only establish it if there are plenty of resources. In that case you could configure the setup priority to be 7 and the holding priority to be 0. In this configuration, the tunnel will never get preempted once established.
Fast Reroute	The value is "True" if Fast Reroute is enabled.
Disjoint Group	If applicable, the disjoint group the RSVP-TE tunnel belongs in.

Column Heading	Description
Disjoint Type	Whether is node, link, or SRLG.
PCE Initiated	The value is "true" if the RSVP-TE tunnel was configured directly on the PCE device. It will be "false" if PCC-initiated.
Actions	Click ··· to get more information about the tunnel. If the tunnel was created using the UI, you can also edit or delete this tunnel.

Visualize SR Policies and RSVP-TE Tunnels

This section describes the visualization features provided in the topology map for TE tunnels that have been discovered during the onboard of devices or provisioned using Cisco Crosswork Optimization Engine. To create and manage TE tunnels using Cisco Crosswork Optimization Engine see Create and Manage SR Policies, on page 107 and Create and Manage RSVP-TE Tunnels, on page 117.

This section contains the following topics:

- Visualize TE Tunnels Example, on page 98
- Highlight a TE Tunnel on the Map, on page 92
- Show IGP, Delay, and Traffic Engineering Metrics, on page 93

Visualize TE Tunnels Example

Follow the steps in this example to quickly familiarize yourself with a number of TE tunnel visualization features that are available from the topology map.

In this example, we are using the following geographical map with devices and links that have SR policies configured. SR policies are not yet highlighted in the map.

Figure 28: Topology Map Example

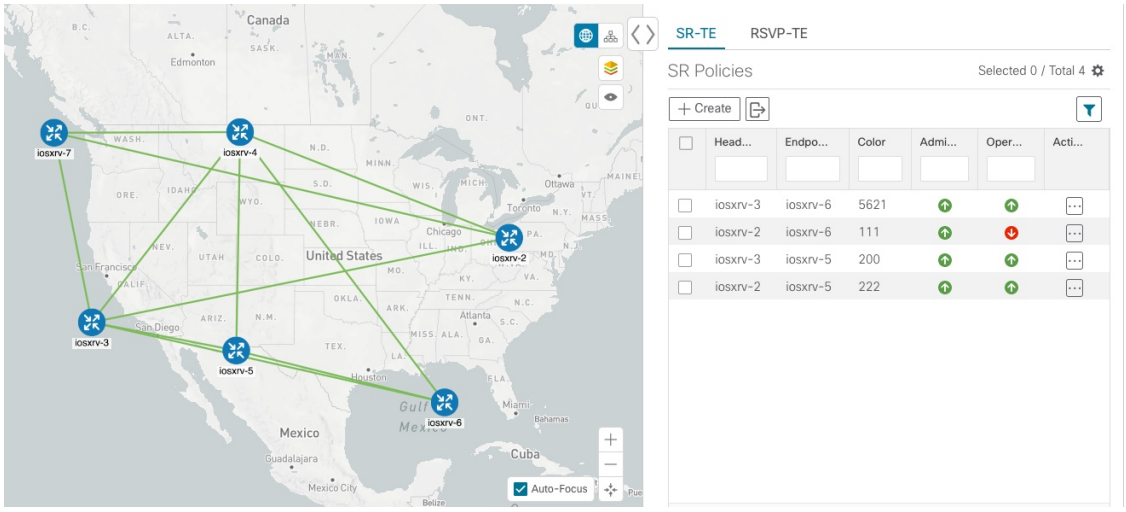

Before you begin

In this example, we assume that devices and SR policies have already been added to Crosswork Optimization Engine (see Get Started, on page 11). While this example uses SR policies, it also applies to RSVP TE tunnels unless otherwise noted.

Step 1 From the **SR Policies** table, click the checkbox next to the SR policies you are interested in. In this example, there are two SR policies selected.

Figure 29: SR Policy Selection

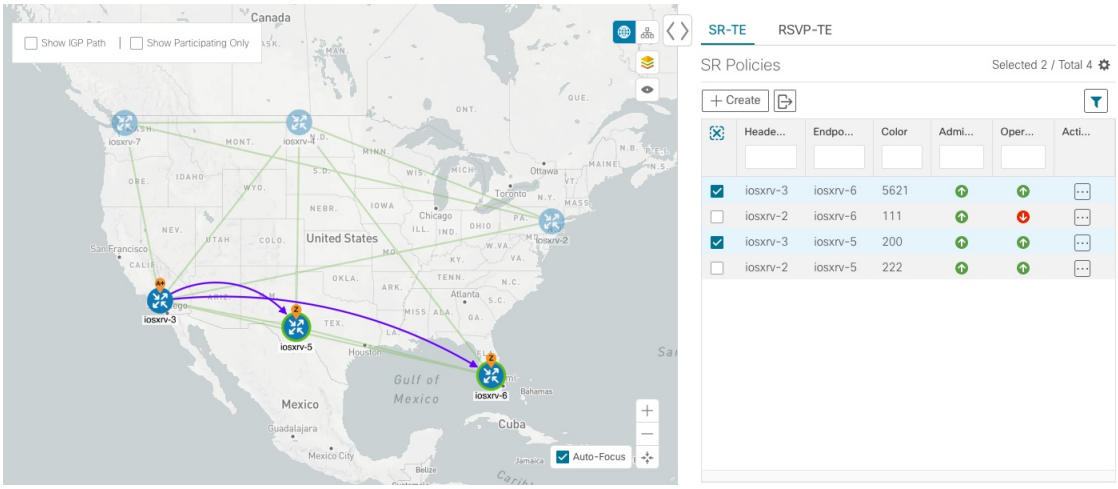

After SR policy selection, the map displays the following:

- SR policies appear as purple links with arrows that indicate the path direction.
- iosxrv-3 is an origin for the both selected policies. iosxrv-5 and iosxrv-6 are destinations for the selected policies. SR policy origin and destination are marked with **A** and **Z**, respectively. The **A+** denotes that there is more than one policy that originates from a device. A Z+ would denote that the device is a destination for more than one policy.

Note If both **A** and **Z** are displayed in a device cluster, at least one device in the cluster is a source and another is a destination.

-  indicates that iosxrv-5 and iosxrv-6 have node SIDs.

Step 2 From the **SR Policies** table, *hover* over a selected policy. The path name of that policy is highlighted on the topology view. You will also see prefix SID information.

Figure 30: Hover over an SR Policy

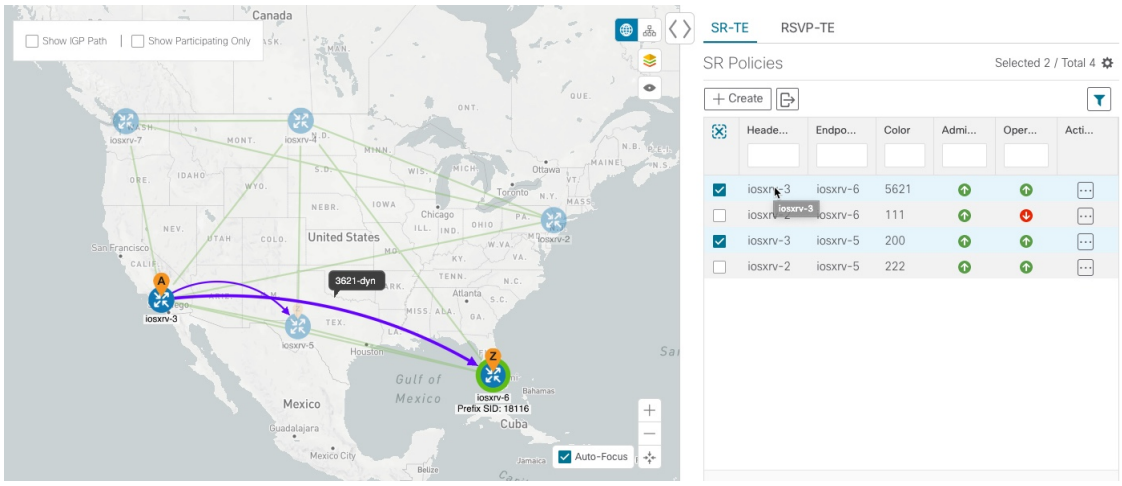

Step 3 Check the **Show IGP Path** check box (available only with SR policies). The IGP paths for the selected SR policies are displayed, with straight lines, instead of the segment hops.

Figure 31: IGP Paths

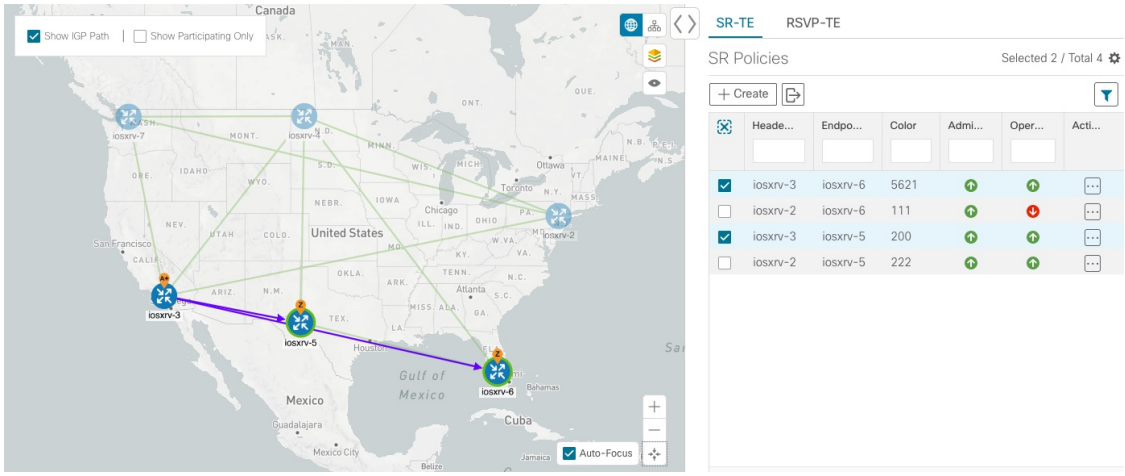

Step 4 Check the **Show Participating Only** check box. All non-participating links and devices disappear. Only participating policies are displayed.

Figure 32: Participating SR Policies

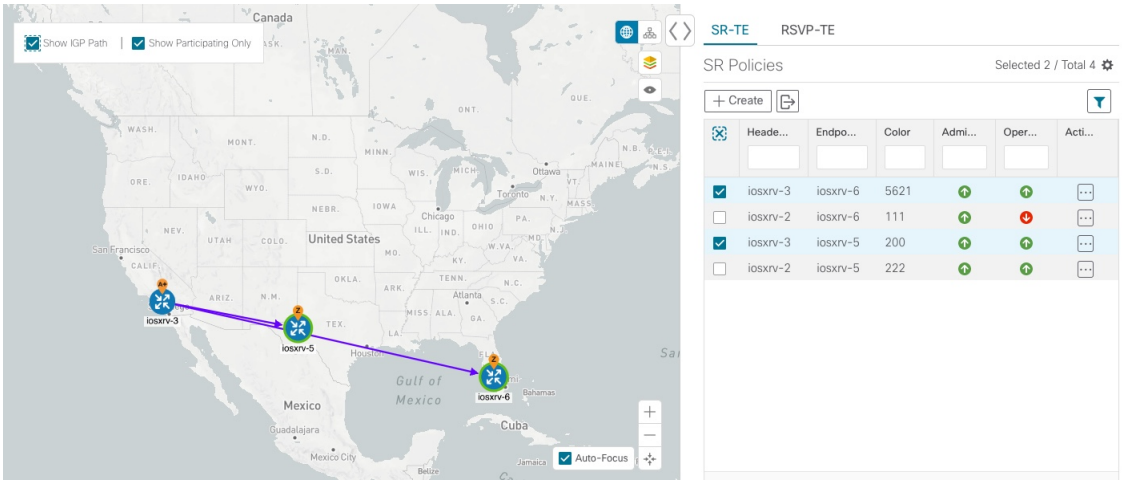

Step 5 To view the IGP, TE or Delay metrics for each tunnel along a policy's path, do the following:

a) For SR policies only, confirm that the Show IGP Path checkbox is checked.

b) Click .

c) Click the **Metrics** tab.

d) Check the applicable metric check boxes.

The metric details are displayed for each policy on the map.

Figure 33: IGP, Delay, and TE Metrics

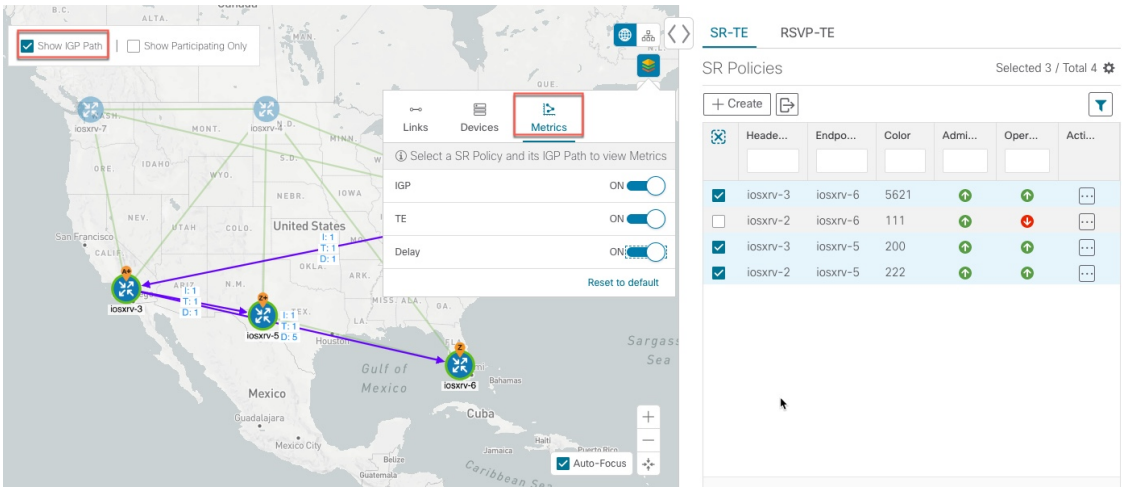

Step 6 Click the logical map icon ( ).

Figure 34: Logical Map

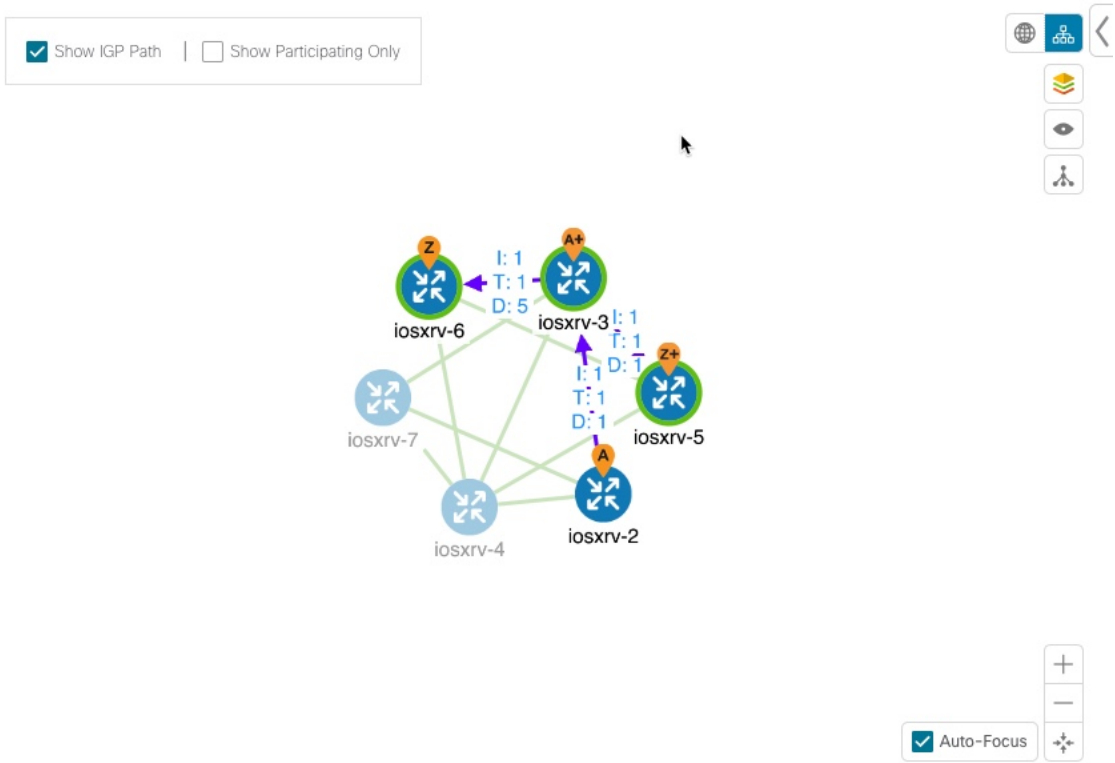

You are able to see the same information (aside from geographical location) that is available on the geographical topology map. You also have the ability to move devices and links on the map to make it easier to view.

Step 7 To view SR policy details such as disjoint groups, metric type, segment hop information, and so on, click ⋯ under the **Actions** column from the table.

The **SR Policy Details** page is displayed in the side panel (see Get More Information About an SR Policy, on page 114). Note that only the selected policy is now highlighted on the topology map.

Figure 35: SR Policy Details

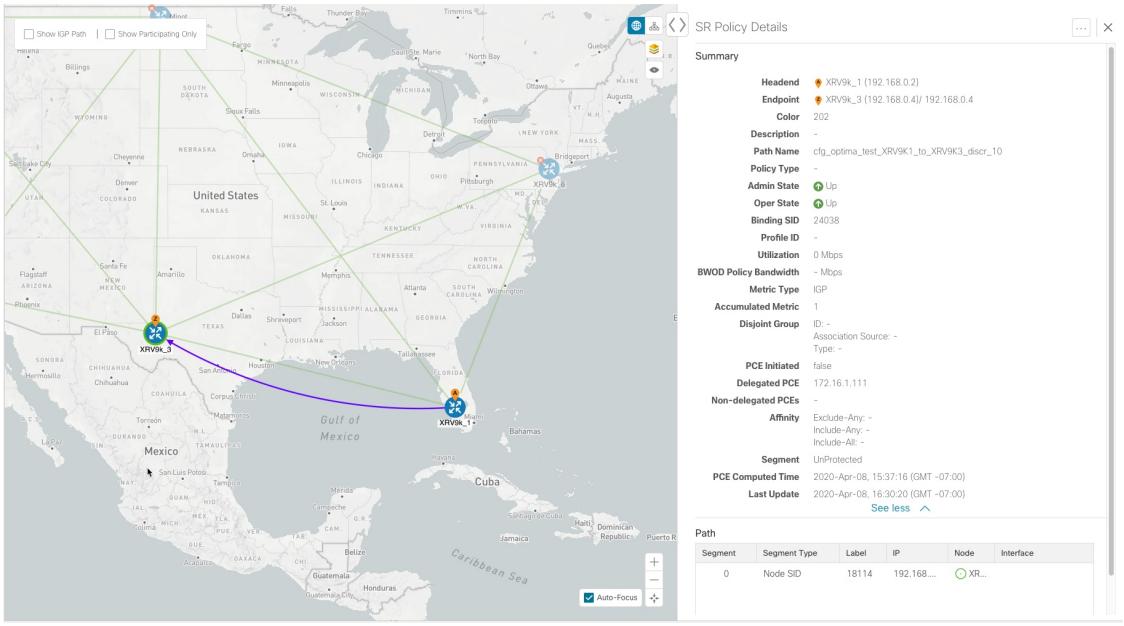

Note To return to the **SR Policies** table, close (✕) the current view.

What to do next

Provision and manage TE tunnels. See Create and Manage SR Policies, on page 107 and Create and Manage RSVP-TE Tunnels, on page 117.

Configure Affinity Mapping

Affinity of a an SR policy or RSVP-TE tunnel is used to specify the link attributes for which the SR policy or RSVP-TE tunnel has affinity for. It determines which links are suitable to form a path for the SR policy or RSVP-TE tunnel. It is a 32-bit value, with each bit position (0 - 31) representing a link attribute. Affinity mapping is used to map each bit position or attribute to a color. This makes it easier to refer to link attributes.

Note The affinity mapping name is only used for visualization in Cisco Crosswork Optimization Engine. Affinities defined on devices are not collected by Cisco Crosswork Optimization Engine. Define affinity mapping in Cisco Crosswork Optimization Engine with the same name and bits that are used on the device interface. Cisco Crosswork Optimization Enginewill only send bit information to SR-PCE during provisioning.

Step 1 From the main menu choose **Optimization Engine** > **Affinity Mapping**. You can also define affinities while creating an SR policy or RSVP-TE tunnel (Create Dynamic Path SR Policies, on page 110 or Create Dynamic Path RSVP-TE Tunnels, on page 118) by clicking **Manage Mapping**.

Step 2 To add a new affinity mapping, click **Create Mapping**.

a) Enter the name (color) and the bit it will be assigned to.

b) Click to save the mapping.

Step 3 To edit an affinity mapping, click .

a) Make the necessary changes. If you want to cancel your changes, click .

b) Click to save the changes.

Step 4 To delete an affinity mapping, click

Note You should remove the TE tunnel before removing the affinity to avoid orphan TE tunnels. If you have removed an affinity associated to a TE tunnel, the affinity is shown as "UNKNOWN" in the **SR Policy / RSVP-TE Tunnel Details** window.

What to do next

After defining affinities, you can Create Dynamic Path SR Policies, on page 110 or Create Dynamic Path RSVP-TE Tunnels, on page 118.

Preview Disjoint SR Policies and RSVP-TE Tunnels

The following example shows how the SR policy and RSVP-TE tunnel provisioning preview feature can be used for disjoint SR policies and RSVP-TE tunnels. In this example, two SR policies will be provisioned with link disjointness. After the first one is provisioned, the preview of the second will show both policies in the map view and how the path of the first would be re-optimized by SR-PCE to make them link disjoint from each other.

Note There cannot be more than 2 disjoint policies in the same disjoint group or subgroup

Below is a provisioned dynamic policy (DisjA) belonging to disjoint link group 200. The SR policy has a path that ECMP splits between XRV9k_4 and XRV9k_1 as shown in the following figure.

Figure 36: Example: DisjA SR Policy

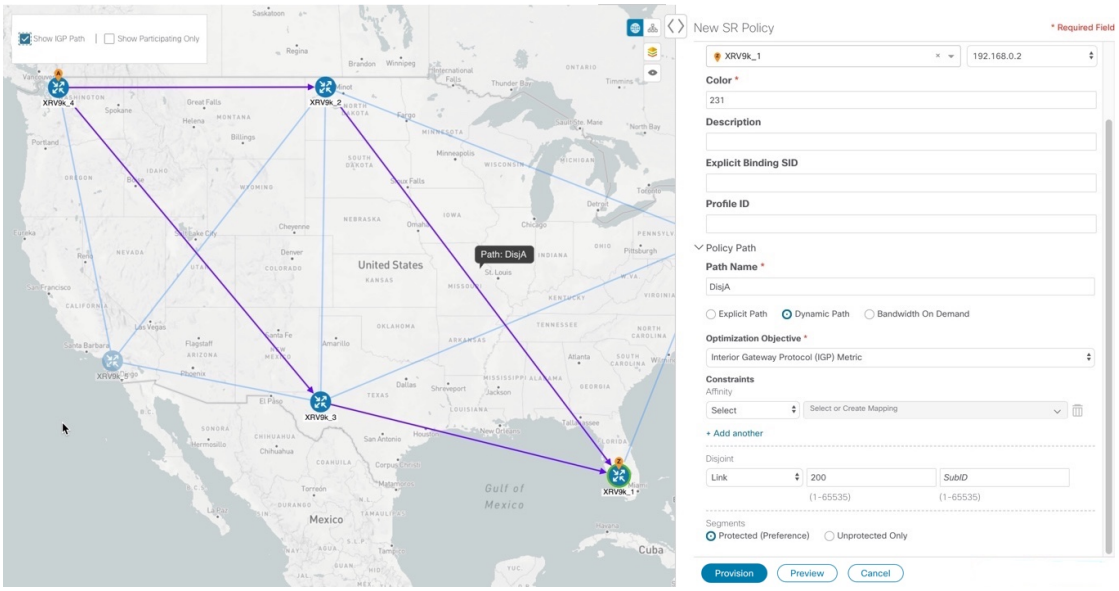

A second policy (DisjB) is now configured in the same disjoint group as the first. When we preview this policy you see both DisjA and DisjB are displayed. You also see the path of DisjA has been reoptimized to ensure both policies are link disjoint. This path change to the existing policy DisjA will be made by SR-PCE if DisjB is provisioned.

Figure 37: Example: Preview Disjoint SR Policies

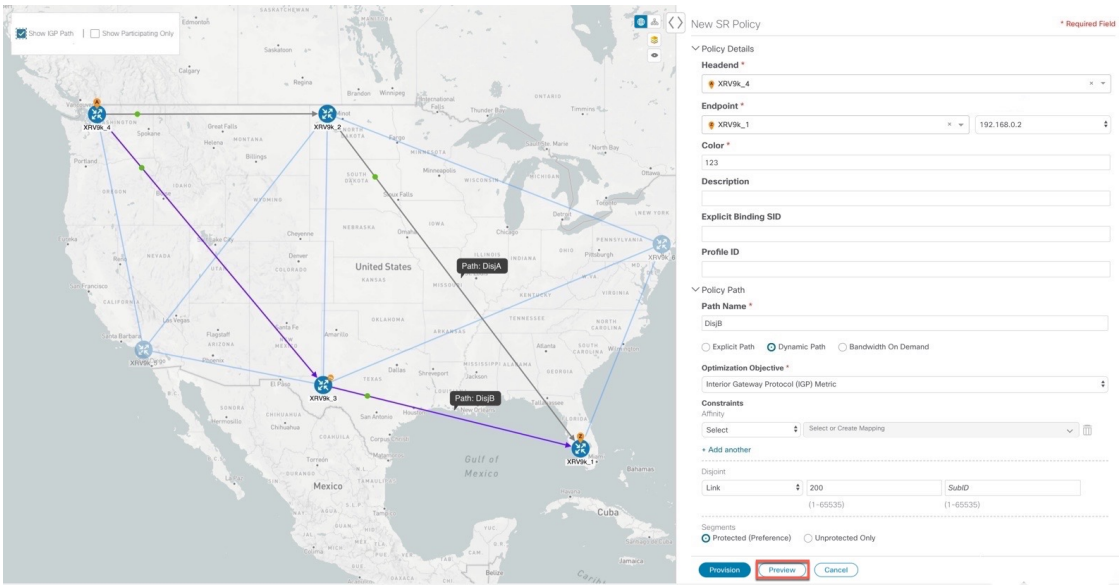

After DisjB is provisioned, we select **View SR Policy List** and check the checkbox next to the DisjA policy to confirm that the path for DisjA has been rerouted.

Figure 38: Example: DisjA SR Policy Rerouted

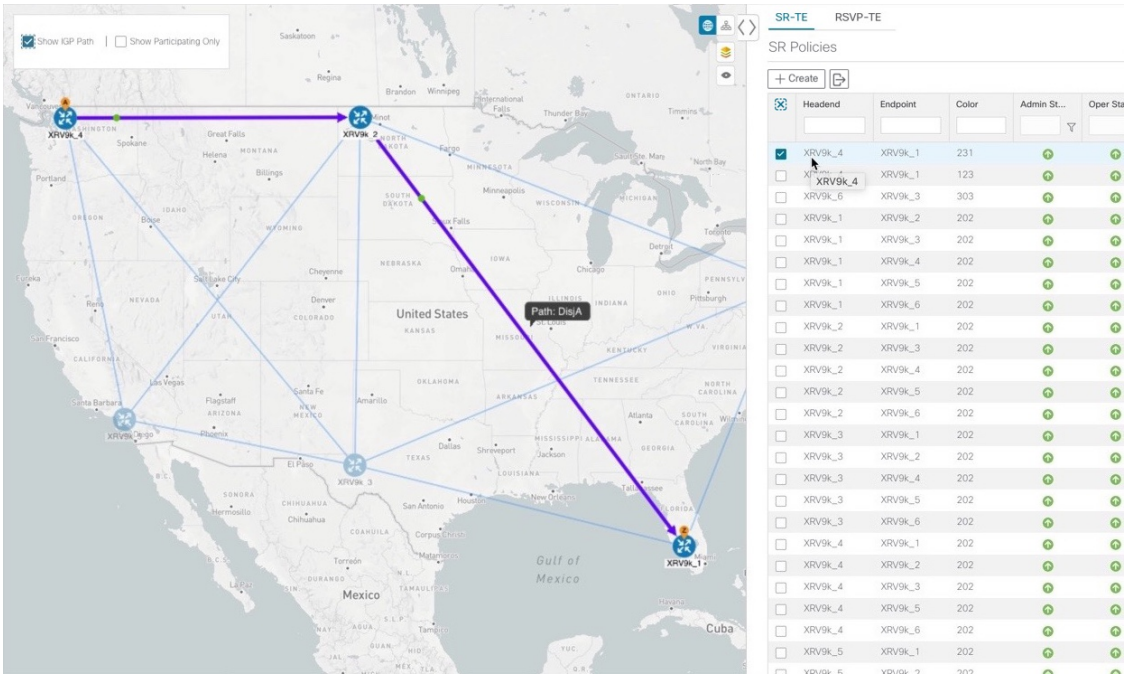

From the SR Policies table, check the checkbox next to DisjB, and delete it.

Figure 39: Example: Delete DisjB SR Policy

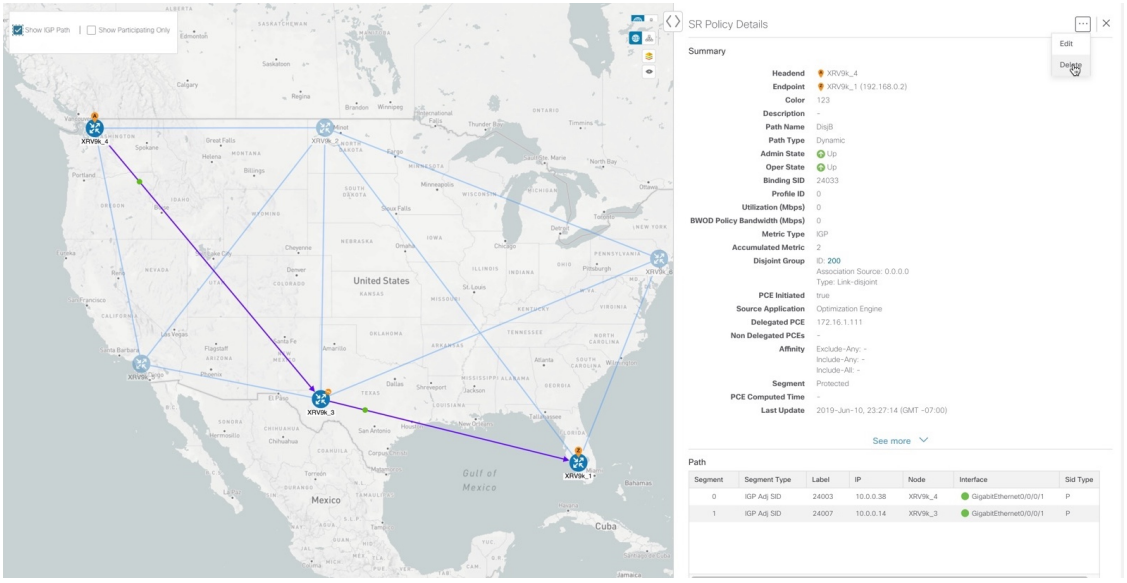

After a few seconds, display DisjA again. You will see that it has reset itself and shows two paths from XR.

Figure 40: Example: DisjA SR Policy Reset

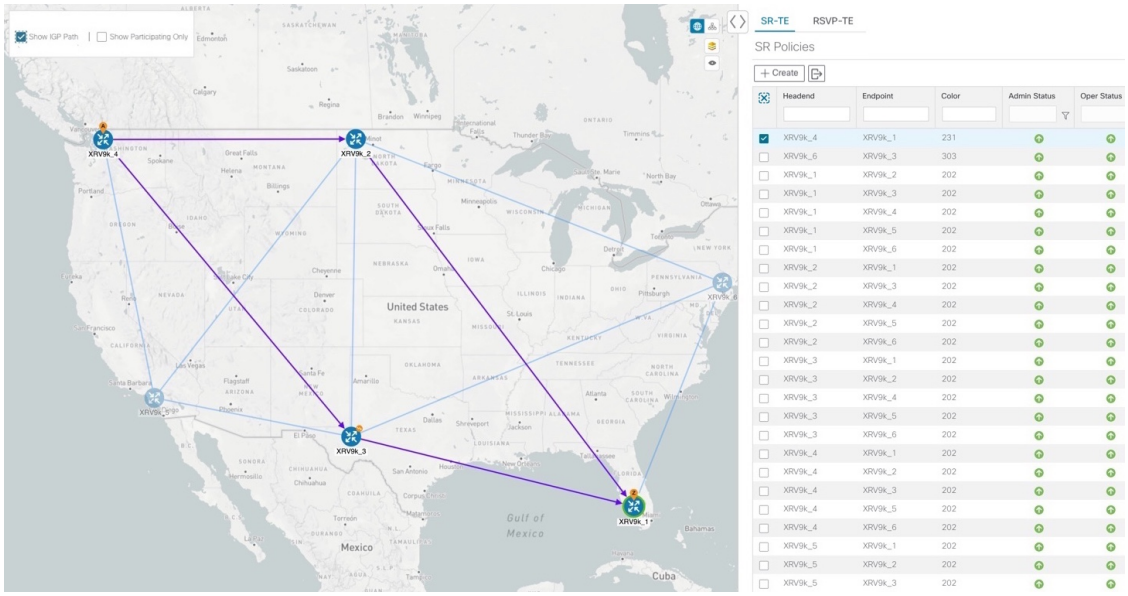

View TE Tunnels Belonging to a Disjoint Group

From the **SR Policy Details** or **RSVP-TE Tunnel Details** window, click the **Disjoint Group** ID number to view all TE tunnels that belongs to the disjoint group.

Figure 41: Disjoint Group

To go back to the **SR Policy Details** or **RSVP-TE Tunnel Details** window, click ↩.

Create and Manage SR Policies

This section describes how to provision and manage SR policies using the Cisco Crosswork Optimization Engine UI. The Cisco Crosswork Optimization Engine UI gives you the capability of provisioning SR policies in a variety of methods (explicit, dynamic, and bandwidth constraint driven). As you provision an SR policy, you can select nodes on the topology map and also preview the path before deployment. This greatly reduces the complexity of SR policy management. Before provisioning SR policies, you should understand some basic segment routing configuration concepts (see Segment Routing, on page 11).

Create Explicit Path SR Policies

This task creates an SR policy using an explicit path (segments) that you define.

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering**.

Step 2 From the **SR Policies** table, click **+ Create**.

Figure 42: Create SR Policy

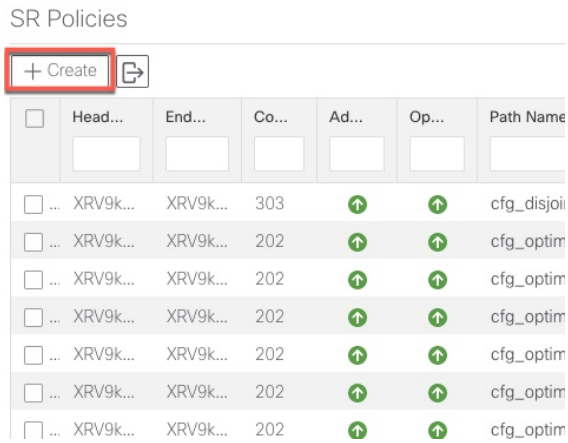

Step 3 Enter the following SR policy values:

a) Required fields:

- **Headend**—Where the SR policy is instantiated. Note: You can either select a node (from the map or drop-down list) or enter part of the node name to filter the headend and endpoint node entries.
- **Endpoint**—The destination of the SR policy.
- **Node Prefix**—After the endpoint is selected, the Node Prefix list is populated and you can select the loopback IP address.
- **Color**—A numerical value that distinguishes between two or more policies to the same node pairs (Headend – Endpoint). Every SR policy between a given headed and endpoint must have a unique color. The bit value must match the value that is configured on the device.
- **Path Name**—Enter a name for this SR policy path. SR policy paths from the same headend must be unique. Policy path names are not case sensitive.

b) Optional values:

- **Description**—Enter details or a description of this policy.
- **Explicit Binding SID**—The binding segment is a local segment identifying an SR policy. Each SR policy is associated with a binding segment ID (BSID). The BSID is a local label that is automatically allocated for each SR policy when the policy is instantiated. If you wish to use a specific segment ID, rather than the default one that is automatically assigned, then enter it here.
- **Profile ID**—Identification used to associate an SR policy with a set of features applied to the policy by the headend. It should correspond with a profile configured on the headend.

Step 4 Under Tunnel Path, click **Explicit Path**.

Step 5 Add segments that are part of the SR policy path.

a) You can either select a node from the drop-down list or enter part of the node name to filter the node list. After a node is selected, the **Select SID** drop-down list is populated with associated prefix and adjacency segment IDs.

b) Select a segment ID from the **Select SID** drop-down list. The drop-down list contains all available segments. The segment names indicate the associated node and whether it is a prefix or an adjacency segment. The name also includes whether the segment is protected (P) or unprotected (U).

c) Click **Add**. The segment appears in the table with segment values.

d) Repeat for each segment you want to add to the SR policy path. To reorder the segment hops, click and drag next to the segment hop you want to move.

Note The segments must be in order or the path will not be created.

Figure 43: Explicit SR Policy Example

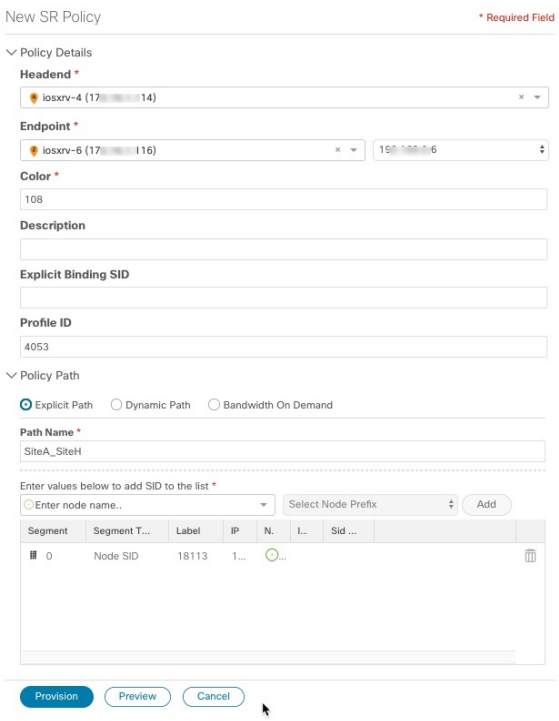

Step 6 Click **Preview**. The path is highlighted on the map and policy details are displayed on the right.

Figure 44: Explicit SR Policy Example

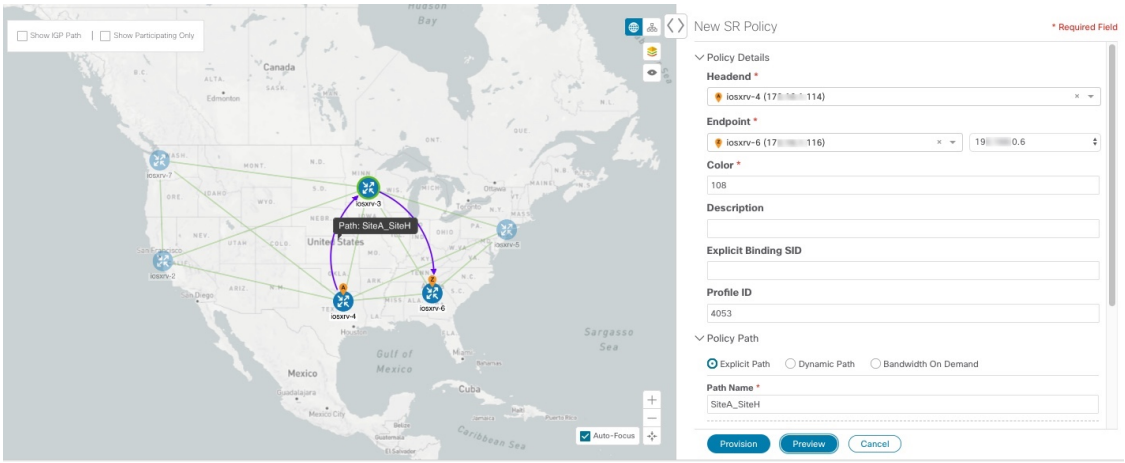

Step 7 If you are satisfied with the policy path, click **Provision**.

Step 8 When the policy is provisioned successfully, a window appears with the following options:

- **View SR Policy List**—Displays the **SR Policies** table that lists all SR policies including the one that was just created.
- **Create New**—Allows you to create another SR policy.

Note The newly provisioned SR policy may take some time, depending on network size and performance, to appear in the **SR Policies** table. The **SR Policies** table is refreshed every 30 seconds.

Create Dynamic Path SR Policies

This task creates an SR policy with a dynamic path. SR-PCE computes a path for the policy based on metrics and path constraints (affinity or disjointness) defined by the user. A user can select from three available metrics to minimize in path computation: IGP, TE, or delay. SR-PCE may also automatically re-optimize the path as necessary based on topology changes.

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering**.

Step 2 From the **SR Policies** table, click **+ Create**.

Figure 45: Create SR Policy

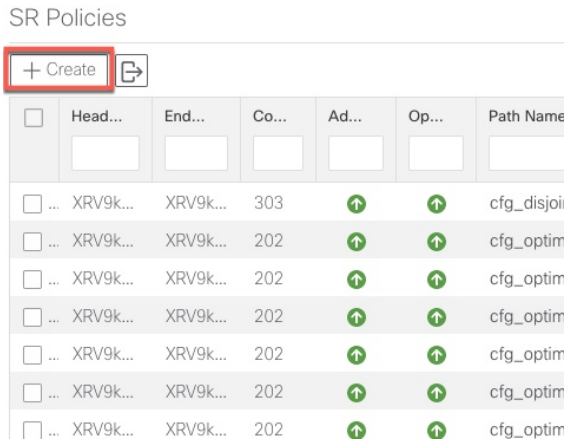

Step 3 Enter the following SR policy values:

a) Required fields:

- **Headend**—Where the SR policy is instantiated. Note: You can either select a node (from the map or drop-down list) or enter part of the node name to filter the headend and endpoint node entries.
- **Endpoint**—The destination of the SR policy.
- **Node Prefix**—After the endpoint is selected, the Node Prefix list is populated and you can select the loopback IP address.
- **Color**—A numerical value that distinguishes between two or more policies to the same node pairs (Headend – Endpoint). Every SR policy between a given headed and endpoint must have a unique color.
- **Path Name**—Enter a name for this SR policy path. SR policy paths from the same headend must be unique. Policy path names are not case sensitive.

b) Optional values:

- **Description**—Enter details or a description of this policy.
- **Explicit Binding SID**—The binding segment is a local segment identifying an SR policy. Each SR policy is associated with a binding segment ID (BSID). The BSID is a local label that is automatically allocated for each SR policy when the policy is instantiated. If you wish to use a specific segment ID, rather than the default one that is automatically assigned, then enter it here.
- **Profile ID**—Identification used to associate an SR policy with a set of features applied to the policy by the headend. It should correspond with a profile configured on the headend.

Step 4 Under Tunnel Path, click **Dynamic Path**.

Step 5 Under Optimization Objective, select one of the following:

- **Interior Gateway Protocol (IGP) Metric**—Minimizes total path IGP metric.
- **Traffic Engineering (TE) Metric**—Minimize total path TE metric.
- **Latency**—Minimize total path latency.

Step 6 Define affinities:

Note Affinity constraints and disjointness cannot be configured on the same SR policy.

- **Exclude Any**—Does not traverse interfaces that have any of the specified affinities.
- **Include Any**—Includes only interfaces that have any of the specified affinities.
- **Include All**—Include only interfaces that have all of the specified affinities.
- **Select or Create Mapping**
  - If affinity mappings have been defined, select the applicable value.
  - To create an affinity mapping, click **Create Mapping**.

  Note For more information, see Configure Affinity Mapping, on page 103.

- **Add Another**—Click this link to add more affinity rules.

Step 7 (Optional) Define disjointness. For more information on how Cisco Crosswork Optimization Engine handles disjoint policies and what options are supported, see the "Disjointness" section in Segment Routing, on page 11). Enter the disjoint group ID and subgroup ID. If there are existing SR policies belonging to a disjoint group that you define here, all SR policies that belong to that same disjoint group are shown during Preview.

Note There cannot be more than two SR policies in the same disjoint group or subgroup.

Step 8 Under Segments, select one of the following:

- **Protected (Preference)**—Creates an SR policy that will use protected segments (provides a backup path) when available.
- **Unprotected Only**—Creates an SR policy that will only use unprotected segments. This option cannot be used when affinity constraints are defined.

Step 9 Click **Preview**. The path is highlighted on the map. Note in the following example that all policies belonging to the same disjoint group are displayed.

Figure 46: Dynamic SR Policy Preview

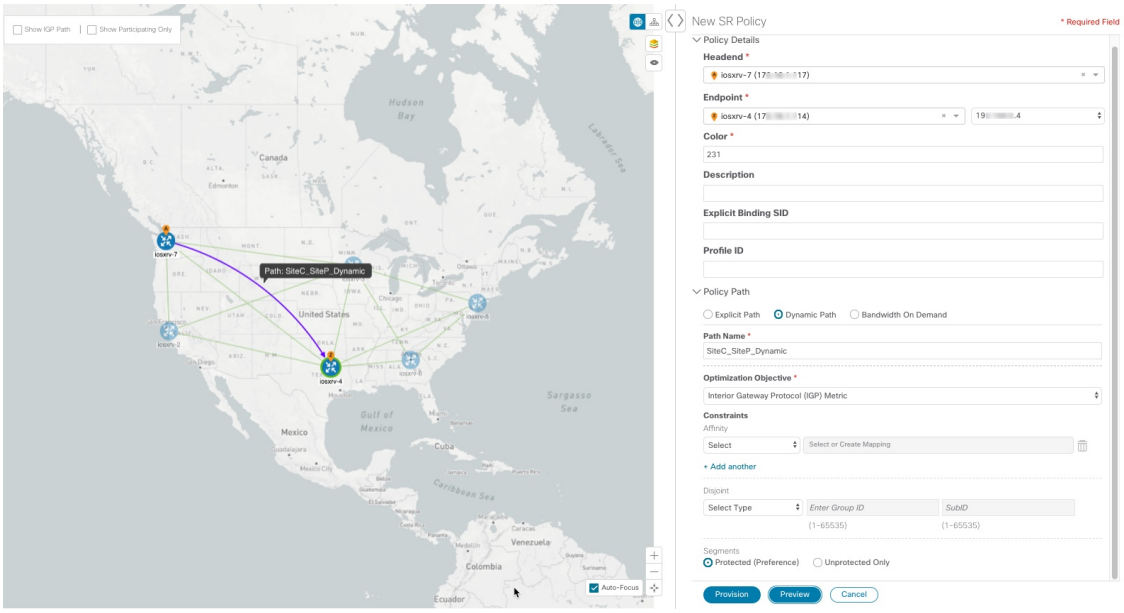

Step 10 If you are satisfied with the policy path, click **Provision**.

Step 11 When the policy is provisioned successfully, a window appears with the following options:

- **View SR Policy List**—Displays the **SR Policies** table that lists all SR policies including the one that was just created.
- **Create New**—Allows you to create another SR policy.

Modify SR Policies

To modify an SR policy:

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering**.

Step 2 Expand the **SR Policies** table. You will see a list of SR policies and various information such as headend, endpoint, Admin status, operating status, and so on.

Step 3 Locate the SR policy you are interested in and click ⋯ (under the **Actions** column). You may need to expand the SR Policies table to view the **Actions** column.

Step 4 From the top-right corner of the **SR Policy Details** window, click ⋯.

Note If the icon is grayed out, the tunnel cannot be modified for one of the following reasons:

- The policy was not created using the Crosswork Optimization Engine UI (**SR Policies** table > **Create**).
- The policy was created using the Bandwidth Optimization function pack.

Step 5 Click **Edit**.

Step 6 In the **Policy Path** area, modify the values you want to change.

Step 7 (Optional) Click **Preview** to view visible updates on the topology map.

Step 8 Click **Update**.

Step 9 When the policy is updated successfully, a window appears with the following options:

- **View SR Policy List**—Displays the **SR Policies** table that lists all SR policies including the one that was just updated.
- **Create New**—Allows you to create a new SR policy.

Delete SR Policies

To delete an SR policy:

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering**.

Step 2 Expand the **SR Policies** table. You will see a list of SR policies and various information such as headend, endpoint, Admin status, operating status, and so on.

Step 3 Locate the policy you are interested in and click ··· (under the **Actions** column). You may need to expand the table to view the **Actions** column.

Step 4 From the top-right corner of the **SR Policy Details** window, click ···.

Note If the icon is grayed out, the tunnel cannot be modified for one of the following reasons:

- The policy was not created using the Crosswork Optimization Engine UI (**SR Policies** table > **Create**).
- The policy was created using the Bandwidth Optimization function pack.

Step 5 Click **Delete**.

Get More Information About an SR Policy

From the **SR Policies** table, locate the SR policy you are interested in and click ··· (under the **Actions** column). You may need to expand the **SR Policies** table to view the **Actions** column. The SR Policy Details window appears. It provides more detailed information about the policy and its associated paths. See the table below for field descriptions.

Figure 47: SR Policy Details

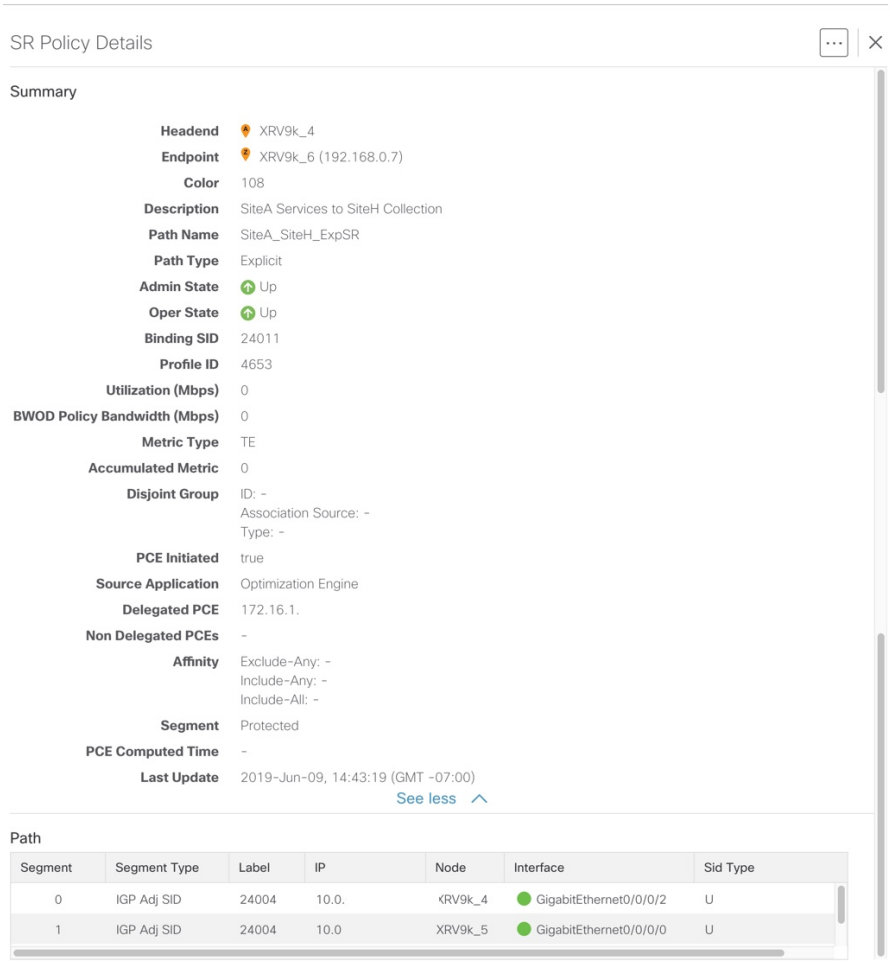

Table 12: SR Policy Details Fields

Field	Description
Headend	Where the SR policy is instantiated (source).
Endpoint	The destination of the SR policy.
Color	A numerical value that distinguishes between two or more policies to the same node pairs (Headend – Endpoint). Every SR policy between a given headend and endpoint must have a unique color.
Description	(Optional) If provisioned using the Cisco Crosswork Optimization Engine UI, it is the description entered by the user. This may be blank if the user did not enter a description.
Path Name	The name of the current active candidate path of the SR policy. For SR policies created using the Cisco Crosswork Optimization Engine UI, it will be the name provided by the user during configuration. For SR policies created through configuration on the headend router, the Path Name will be the base name configured for the policy on the CLI with "cfg_" appended to the beginning and the candidate path preference appended to the end.

Field	Description
Policy Type	Indicates whether an SR policy created through Cisco Crosswork Optimization Engine is explicit or dynamic.
Admin State	Administrative state is dictated by the user. For example, the user creates an SR policy and does not intentionally shut it down. The Admin State will be UP.
Oper State	Operational state received by the system. For example, the user has configured a policy and so the Admin State is UP. However, due to network issues it is operationally down. In this case, Oper State will display DOWN and Admin State will remain as UP.
Binding SID	The binding segment is a local segment identifying an SR policy. Each SR policy is associated with a binding segment ID (BSID). The BSID is a local label that is automatically allocated (or explicitly entered during manual provisioning) for each SR policy when the policy is instantiated.
Profile ID	Identification used to associate an SR policy with a set of features applied to the policy by the headend. It should correspond with a profile configured on the headend.
Utilization (Mbps)	The measured traffic on the SR policy.
BWOD Policy Bandwidth (Mbps)	The bandwidth constraint associated with a policy created through the Bandwidth on Demand function pack.
Metric Type	The metric type can be TE, IGP, or latency.
Accumulated Metric	Total metric calculation of the SR policy.
Disjoint Group	If applicable, displays disjointness information.
PCE Initiated	If the policy was initiated and provisioned by a PCE, the value is **True**.
Delegated PCE	The SR policy is delegated to this PCE IP address.
Non Delegated PCEs	PCEs reporting the policy, but not currently delegated.
Affinity	Lists any affinity constraints belonging to this policy.
Segment	Lists whether a dynamic path policy should prefer protected or require unprotected SIDs
PCE Computed Time	Time when PCE computed the path currently in effect.
Last Update	The last time the policy was updated.
Path	Lists segments that are part of the policy. It gives the following segment information: segment type, label, IP address, associated node, interface, and SID type (Protected or Unprotected).

Create and Manage RSVP-TE Tunnels

This section describes how to provision and manage RSVP-TE Tunnels using the Cisco Crosswork Optimization Engine UI. As you provision an RSVP-TE Tunnel, you can select nodes on the topology map and also preview the path before deployment. This greatly reduces the complexity of RSVP-TE Tunnel management.

Create Explicit Path RSVP-TE Tunnels

This task creates an RSVP-TE tunnel using an explicit path (hops) that you define.

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering** and select the **RSVP-TE** tab.

Step 2 From the **RSVP-TE** table, click **+ Create**.

Step 3 Enter the following RSVP-TE Tunnel values:

a) Required fields (labeled with red asterisk):

- **Headend**—Where the RSVP-TE tunnel is instantiated. Note: You can either select a node (from the map or drop-down list) or enter part of the node name to filter the headend and endpoint node entries.
- **Endpoint**—The destination of the RSVP-TE tunnel.
- **Path Name**—User specified name for the RSVP-TE tunnel.

Optional fields:

- **Description**—Details or a description of this TE tunnel.
- **Binding Label**—Numeric value of the binding label assigned to this tunnel. By default, the system will assign a value if the user does not enter one.
- **Signaled Bandwidth**—Required bandwidth.
- **Setup Priority**—The default value is 7. There are 8 (0 - 7) setup priorities. 0 is the most preferred. The setup priority is used to define preference for preempting less preferred tunnels. The most preferred tunnels can push the other less preferred tunnels out of the way.
- **Hold Priority**—The default value is 7. There are 8 (0 - 7) hold priorities. The holding priority is used to define a priority maintaining the currently established tunnel. You can have a tunnel that you never want go down, but only establish it if there are plenty of resources. In that case you could configure the setup priority to be 7 and the holding priority to be 0. In this configuration, the tunnel will never get preempted once established.
- **Fast Reroute**—By default, Fast Re-route (FFR) is disabled. FFR provides fast traffic recovery when links fail.

Step 4 Under Tunnel Path, click **Explicit Path**.

Step 5 Add hops that will be part of the RSVP-TE tunnel.

a) Select a node from the drop-down node list.

b) Select the IP address from the **Select Interface** drop-down list. The drop-down list contains all available hops.

c) Select **Strict** as the type of hop. A strict path means that a network node and its preceding node in the ERO must be adjacent and directly connected. Each strict hop should be specified as a remote (ingress) interface or the Loopback IP (TE router ID) of the node.Crosswork Optimization Engine does not support configuration of loose hops in this release.

d) Click **Add**.

e) Repeat for each hop you want to add to the RSVP-TE tunnel. To reorder the hops, click and drag ⠿ next to the hop you want to move.

Note The hops must be in order or the path will not be created.

Step 6 Click **Preview**. The path is highlighted on the map and policy details are displayed on the right. See the following figure to see a sample of RSVP-TE tunnel configuration details and a preview of the new tunnel.

Figure 48: Explicit RSVP-TE Tunnel Example

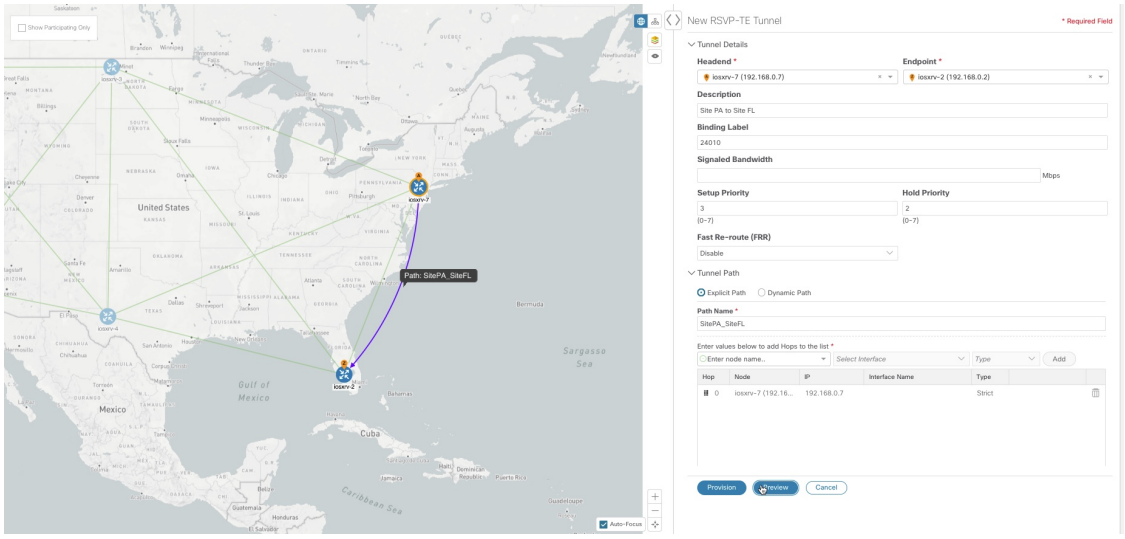

Step 7 If you are satisfied with the path, click **Provision**.

Step 8 When the RSVP-TE tunnel is provisioned successfully, a window appears with the following options:

- **View RSVP-TE List**—Displays the **RSVP-TE Tunnels** table that lists all RSVP-TE tunnels including the one that was just created.
- **Create New**—Allows you to create another TE tunnel.

Create Dynamic Path RSVP-TE Tunnels

This task creates an RSVP-TE tunnel with a dynamic path. SR-PCE computes a path for the tunnel based on metrics and path constraints (affinity or disjointness) defined by the user. A user can select from three available metrics to minimize in path computation: IGP, TE, or delay. SR-PCE may also automatically re-optimize the path as necessary based on topology changes.

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering**.

Step 2 From the **RSVP-TE Tunnel** table, click **+ Create**.

Step 3 Enter the following RSVP-TE Tunnel values:

a) Required fields (labeled with red asterisk):

- **Headend**—Where the RSVP-TE tunnel is instantiated. Note: You can either select a node (from the map or drop-down list) or enter part of the node name to filter the headend and endpoint node entries.
- **Endpoint**—The destination of the RSVP-TE tunnel.
- **Path Name**—User specified name for the RSVP-TE tunnel.

Optional fields:

- **Description**—Details or a description of this TE tunnel.
- **Binding Label**—Numeric value of the binding label assigned to this tunnel. By default, the system will assign a value if the user does not enter one.
- **Signaled Bandwidth**—Required bandwidth.
- **Setup Priority**—The default value is 7. There are 8 (0 - 7) setup priorities. 0 is the most preferred. The setup priority is used to define preference for preempting less preferred tunnels. The most preferred tunnels can push the other less preferred tunnels out of the way.
- **Hold Priority**—The default value is 7. There are 8 (0 - 7) hold priorities. The holding priority is used to define a priority maintaining the currently established tunnel. You can have a tunnel that you never want go down, but only establish it if there are plenty of resources. In that case you could configure the setup priority to be 7 and the holding priority to be 0. In this configuration, the tunnel will never get preempted once established.
- **Fast Reroute**—By default, Fast Re-route (FFR) is disabled. FFR provides fast traffic recovery when links fail.

Step 4 Under Tunnel Path, click **Dynamic Path**.

Step 5 Under Optimization Objective, select one of the following:

- **Interior Gateway Protocol (IGP) Metric**—Minimizes total path IGP metric.
- **Traffic Engineering (TE) Metric**—Minimize total path TE metric.
- **Latency**—Minimize total path latency.

Step 6 Define affinities:

Note Affinity constraints and disjointness cannot be configured on the same tunnel.

- **Exclude Any**—Does not traverse interfaces that have any of the specified affinities.
- **Include Any**—Includes only interfaces that have any of the specified affinities.
- **Include All**—Include only interfaces that have all of the specified affinities.
- **Select or Create Mapping**
  - If affinity mappings have been defined, select the applicable value.
  - To create an affinity mapping, click **Create Mapping**.

  Note For more information, see Configure Affinity Mapping, on page 103.

- **Add Another**—Click this link to add more affinity rules.

Step 7 (Optional) Define disjointness. For more information on how Cisco Crosswork Optimization Engine handles disjoint tunnels and what options are supported, see the "Disjointness" section in Segment Routing, on page 11 (applies to RSVP-TE tunnels as well). Enter the disjoint group ID and subgroup ID. If there are existing tunnels belonging to a disjoint group that you define here, all tunnels that belong to that same disjoint group are shown during Preview.

Note There cannot be more than two TE tunnels in the same disjoint group or subgroup.

Step 8 Click **Preview**. The path is highlighted on the map. See the following figure to see a sample of RSVP-TE tunnel configuration details and a preview of the new tunnel.

Figure 49: Dynamic RSVP-TE Tunnel Preview

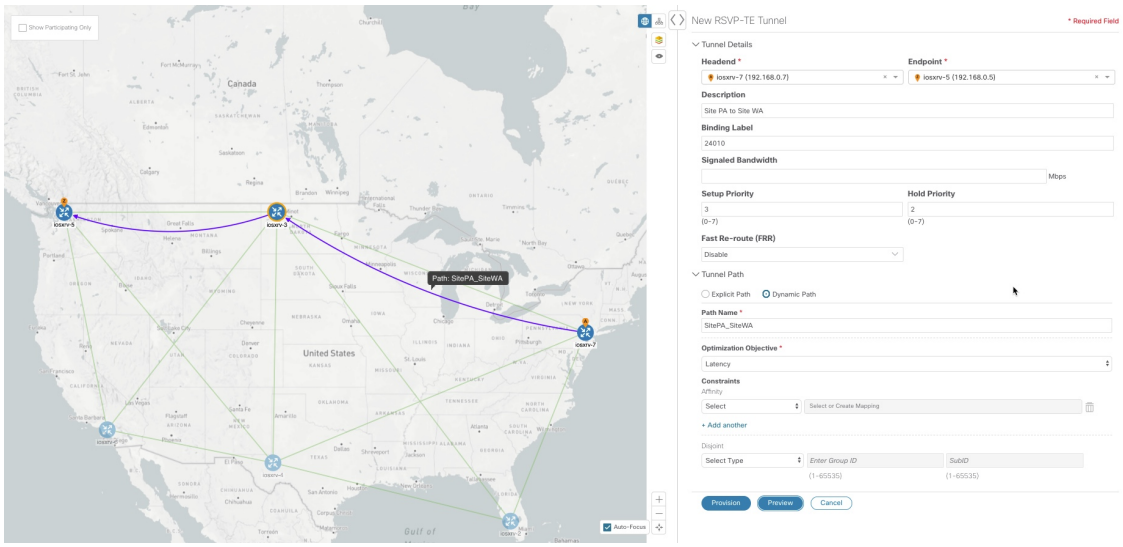

Step 9 If you are satisfied with the tunnel path, click **Provision**.

Step 10 When the tunnel is provisioned successfully, a window appears with the following options:

- **View RSVP List**—Displays the **RSVP-TE Tunnels** table that lists all RSVP-TE tunnels including the one that was just created.
- **Create New**—Allows you to create a new RSVP-TE tunnel.

Modify RSVP-TE Tunnels

To modify an RSVP-TE tunnel:

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering**.

Step 2 Expand the **RSVP-TE Tunnels** table. You will see a list of RSVP-TE tunnels and various information such as headend, endpoint, Admin status, operating status, and so on.

Step 3 Locate the RSVP-TE tunnel you are interested in and click ⋯ (under the **Actions** column). You may need to expand the table to view the **Actions** column.

Step 4 From the top-right corner of the **RSVP-TE Tunnel Details** window, click ⋯.

Note If the icon is grayed out, the policy cannot be modified because the policy was not created using the Crosswork Optimization Engine UI (**RSVP-TE Tunnel** table > **+ Create** button).

Step 5 Click **Edit**.

Step 6 Modify the values you want to change.

Step 7 (Optional) Click **Preview** to view visible updates on the topology map.

Step 8 Click **Update**.

Step 9 When the tunnel is updated successfully, a window appears with the following options:

- **View RSVP List**—Displays the **RSVP-TE Tunnels** table that lists all RSVP-TE tunnels including the one that was just updated.
- **Create New**—Allows you to create a new RSVP-TE tunnel.

Delete RSVP-TE Tunnels

To delete an RSVP-TE tunnel:

Step 1 From the main menu, choose **Optimization Engine** > **Traffic Engineering**.

Step 2 Expand the **RSVP-TE Tunnels** table. You will see a list of RSVP-TE tunnels and various information such as headend, endpoint, Admin status, operating status, and so on.

Step 3 Locate the RSVP-TE tunnel you are interested in and click ⋯ (under the **Actions** column). You may need to expand the table to view the **Actions** column.

Step 4 From the top-right corner of the **RSVP-TE Tunnel Details** window, click ⋯.

Note If the icon is grayed out, the policy cannot be deleted because the tunnel was not created using the Crosswork Optimization Engine UI (**RSVP-TE Tunnel** table > **+ Create** button) or if it was created from another Crosswork Optimization Engine VM that knows of the same topology.

Step 5 Click **Delete**.

Get More Information About an RSVP-TE Tunnel

From the **RSVP-TE Tunnel** table, locate the TE tunnel you are interested in and click the ⋯ link (under the **Actions** column). You may need to expand the **RSVP-TE Tunnel** table to view the **Actions** column. The RSVP-TE Tunnel Details window appears, where you can view more detailed information about the TE tunnel and its associated paths. See the following table for field descriptions.

Figure 50: RSVP-TE Tunnel Details

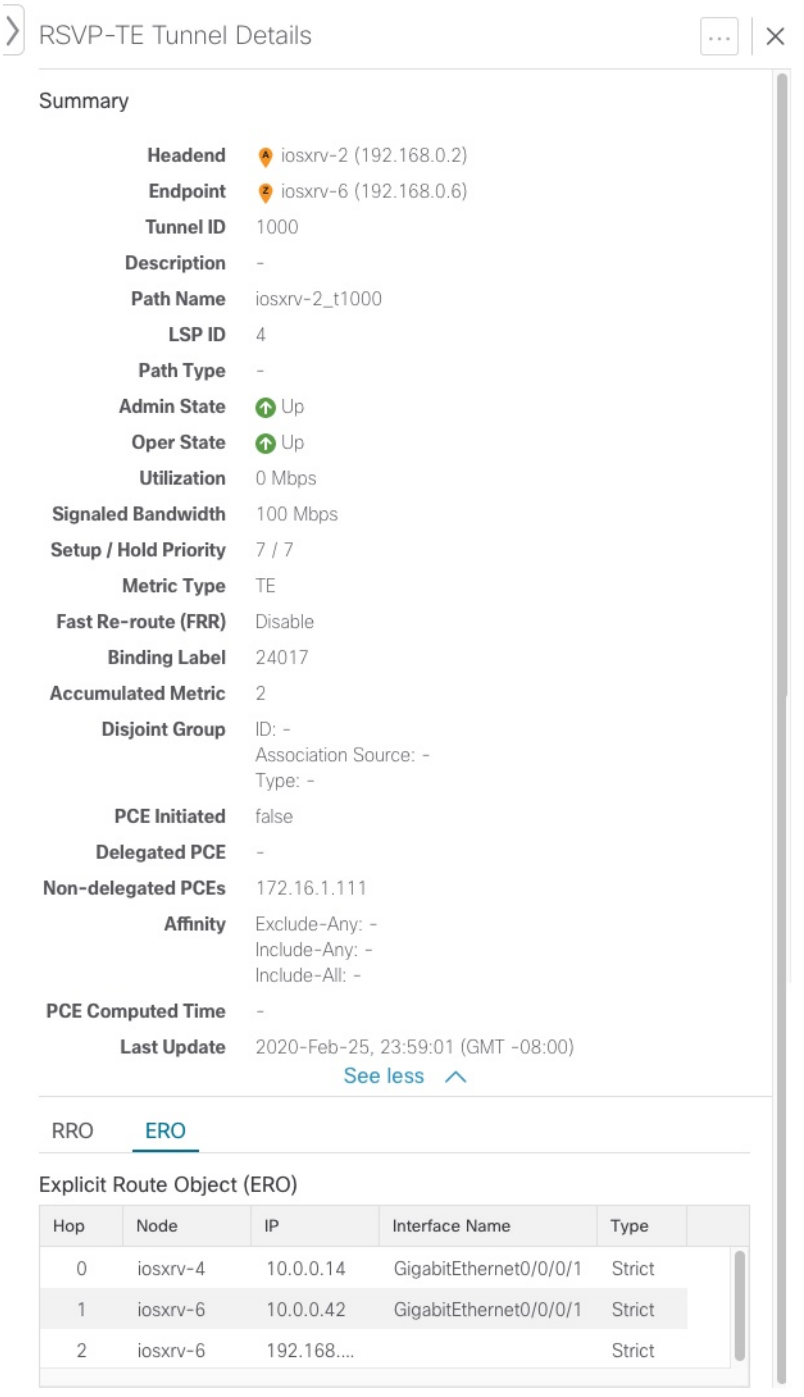

Table 13: RSVP-TE Tunnels

Field	Description
Headend	Where the RSVP-TE tunnel is instantiated.
Endpoint	The destination of the RSVP-TE tunnel.
Tunnel ID	Assigned RSVP-TE tunnel ID.
Path Name	For RSVP-TE tunnels created using the Cisco Crosswork Optimization Engine UI, it will be the name provided by the user during configuration. For RSVP-TE tunnels created through configuration on the headend router, the Path Name for Cisco PCCs will be an auto-generated string consisting of the node name as well as the Tunnel ID.
LSP ID	The LSP identification number.
Path Type	Indicates whether the TE tunnel created through Cisco Crosswork Optimization Engine is explicit or dynamic.
Admin State	Administrative status of the RSVP-TE tunnel. This is the status defined by the user.
Oper State	Operational status of the RSVP-TE tunnel. This is the state of the policy as reported by the system. For example, the user can define the Admin status as Up. However, if the policy is operationally down due to some network issues, then the Oper Status will display as Down.
Utilization	Tunnel's utilization against bandwidth.
Signaled Bandwidth	Bandwidth requirements.
Setup / Hold Priority	There are 8 (0 - 7) setup priorities. 0 is the most preferred. The setup priority is used to define preference for preempting less preferred tunnels. The most preferred tunnels can push the other less preferred tunnels out of the way. There are 8 (0 - 7) hold priorities. The holding priority is used to define a priority maintaining the currently established tunnel. You can have a tunnel that you never want go down, but only establish it if there are plenty of resources. In that case you could configure the setup priority to be 7 and the holding priority to be 0. In this configuration, the tunnel will never get preempted once established.
Metric Type	Type of metric (IGP, TE, or Delay).

Field	Description
Fast Re-route (FRR)	The value is **Enable** if Fast Reroute is enabled.
Binding Label	Defined binding SID label.
Accumulated Metric	Total metric calculation of the RSVP-TE tunnel.
Disjoint Group	If applicable, the disjoint group details the RSVP-TE tunnel belongs in.
PCE Initiated	If the RSVP-TE tunnel was initiated and provisioned by a PCE, the value is **true**. If it is PCC-initiated, the value is **false**.
Delegated PCE	If applicable, the RSVP-TE tunnel is delegated to this PCE IP address.
Non-delegated PCE	PCEs reporting the RSVP-TE tunnel, but not currently delegated.
Affinity	Lists any affinity constraints belonging to this TE tunnel.
PCE Computed Time	Time when PCE computed the path currently in effect.
Last Update	The last time the policy was updated.
Explicit Route Object (ERO)	Lists hop EROs that are part of the tunnel. It gives the following information: node, IP address, interface, and type (strict or loose). **Note** When the ERO tab is selected, the topology map displays the paths as curved lines. If both RRO and ERO paths are available, the RRO path is displayed by default.
Record Route Object (RRO)	Lists hop RROs that are part of the tunnel. It gives the following information: node, IP address, and interface. **Note** When the RRO tab is selected, the topology map displays the paths as straight lines. If both RRO and ERO paths are available, the RRO path is displayed by default.

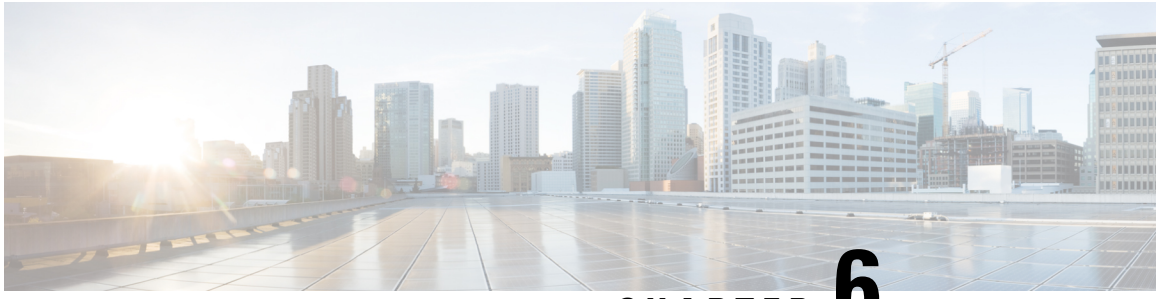

CHAPTER 6

Perform Administrative Tasks

This section contains the following topics:

- Manage Cisco Crosswork Network Automation, on page 125
- Manage Backup and Restore, on page 134
- Manage Users, on page 137
- Manage TACACS+ Servers, on page 142
- Manage LDAP Servers, on page 143
- Define Network Topology Display Settings, on page 145
- Manage Certificates, on page 145
- Smart Licensing Registration, on page 148
- Security Hardening Overview, on page 154

Manage Cisco Crosswork Network Automation

The **Crosswork Manager** window gives you consolidated information about the current status of each installed Cisco Crosswork Optimization Engine application and its supporting services. It also supplies tools and information that, with support and guidance from your Cisco Customer Experience account team, you can use to identify, diagnose and fix issues with Cisco Crosswork Optimization Engine.

Select **Admin** > **Crosswork Manager** to display a **Crosswork Manager** window, with information like the window shown in the following example.

Figure 51: Crosswork Manager Window

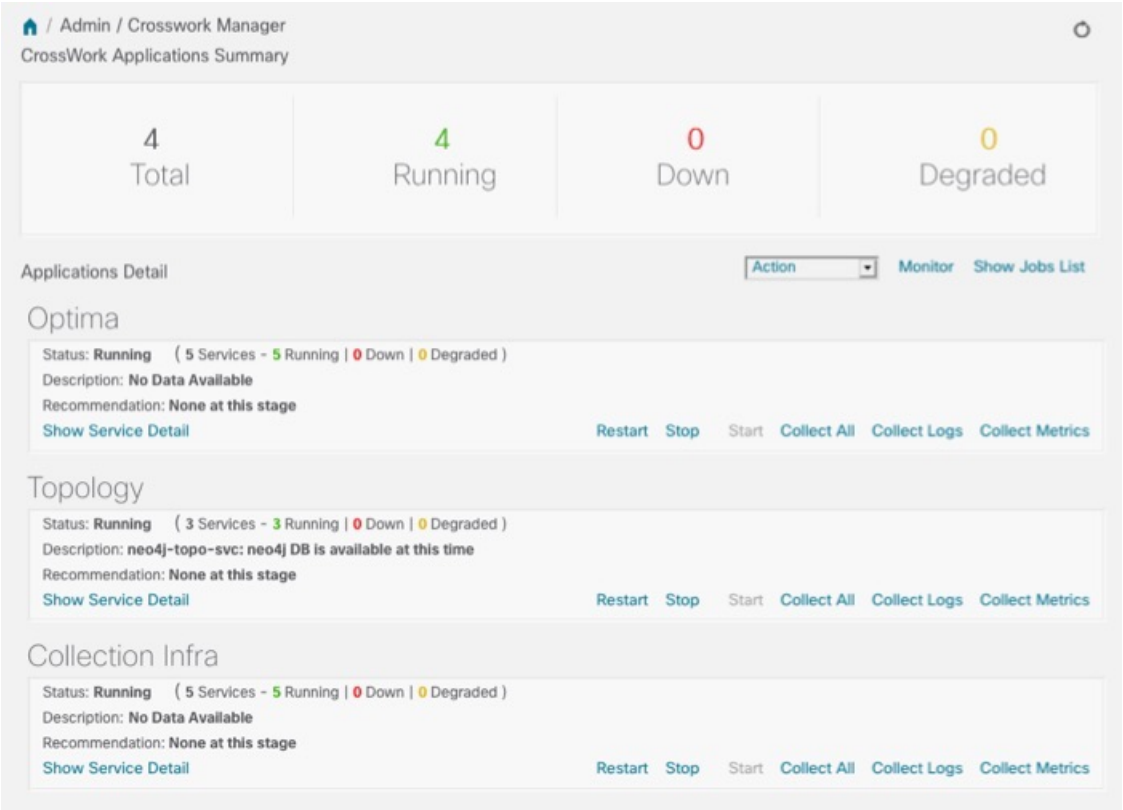

The **Crosswork Manager** window has two main views. The **Crosswork Applications Summary** view, at the top of the window, is a dashboard giving you a quick look at the overall health of the system. It displays the total number of Cisco Crosswork Optimization Engine applications currently installed in the system, and how many of that total are **Running**, **Down**, or **Degraded**.

The **Applications Detail** view, below the **Crosswork Applications Summary** view, allows you to:

- View the name and current runtime status of each installed application and its supporting services.
- Get advice about what to do when an application or one of its services has issues.
- Collect logs and metrics on any application or service, or for the system as a whole.
- Stop, start, or restart any application or service.

The **Applications Detail** view, shown in the following figure, is the best way to investigate any system health issues indicated in the **Crosswork Applications Summary**.

Figure 52: Applications Detail View

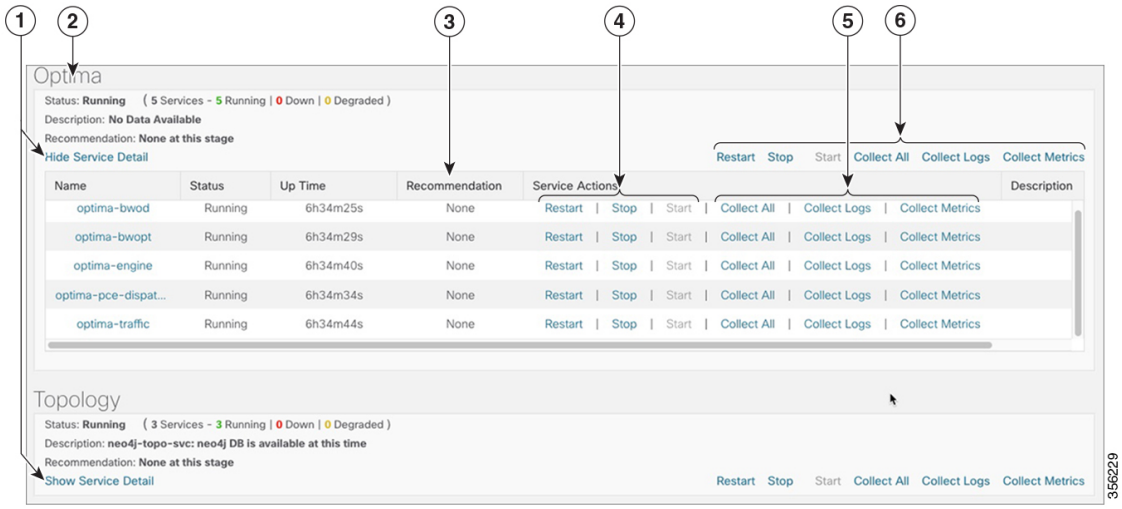

Figure 53: Applications Detail View

Item	Description
1	Click the **Show/Hide Service Detail** link in each application tile to view the detailed status of the underlying services for that application.
2	An **application tile** like this shows the current status of the named application and a summary of the status of that application's services. This includes the total number of services, and how many of those services are Running, Down, or Degraded.
3	Both the **application tile** and its **Service Detail** table provide the name, status, description and recommendation for the respective application or service. The Service Detail table also provides service uptime, and you can click on the link in the **Name** column to see more details about the service, such as its process ID and pod identifier.
4	To control an application or service, click on any of the links in this section of the application tile or Service Detail table. You can click: • **Restart** to restart the application or service. • **Stop** to stop the application or service. • **Start** to start the application or service. See Control Cisco Crosswork Network Automation Applications and Services, on page 133.

Item	Description
5	To gather logs and metrics for the entire system, or for any application or service, click on any of the "collect" links at the system (in the dropdown menu), application, or service level. You can choose: • **Collect All** to collect both logs and metrics. • **Collect Logs** to collect only logs. • **Collect Metrics** to collect only metrics. See Collect and Share Cisco Crosswork Network Automation Logs and Metrics, on page 132.
6	Click the **Monitor** link to monitor individual Cisco Crosswork Optimization Engine functions and features, using analytical dashboards and data gathered over the last 24 hours of run time. See Monitor Cisco Crosswork Network Automation Functions in Real Time, on page 128.
7	Choosing any of the control or collect actions at the system, application or service level will initiate a job. You can view each job's progress by clicking the **Show Jobs List**link at the top right corner of the window. You can also use the **Show Jobs List**to publish collected logs and metrics files, and check on the status of publish jobs you initiate.

Monitor Cisco Crosswork Network Automation Functions in Real Time

You can monitor the health of Cisco Crosswork Optimization Engine and any of its functions in real time, using a set of monitoring dashboards you can access from the **Crosswork Manager** window.

Cisco Crosswork Optimization Engine uses Grafana to create these dashboards. They give you a graphical view of the product's infrastructure, using metrics collected in its database. You can use these dashboards to diagnose problems you may encounter with individual Cisco Crosswork Optimization Engine applications or their underlying services.

There are multiple monitor dashboards, categorized by the type of functionality they monitor and the metrics they provide, as shown in the following table.

Table 14: Monitoring Dashboard Categories

This dashboard category...	Monitors...
Optima	Cisco Crosswork Optimization Engine function pack, traffic, and SR-PCE dispatcher functions.
Topology	Topology service and database functions.
Collection Infra	Device-data collection functions. Metrics include telemetry collection latencies, total collection operations, memory and database activity related to telemetry, delayed collections, and so on.
Core Infra	System hardware and communications usage and performance. Metrics include disk and CPU usage, database size, network and disk operations, and client/server communications.

To conserve disk space, Cisco Crosswork Optimization Engine maintains a maximum of 24 hours of collected metric data.

Grafana is an open-source visualization tool. The following provides general information about how to use the Cisco Crosswork Optimization Engine implementation of Grafana. For more information about Grafana itself, see https://grafana.com and http://docs.grafana.org

Step 1 From the main menu, choose **Admin** > **Crosswork Manager**.

Step 2 At the right, just below the **Crosswork Applications Summary** view, click the **Monitor** link, highlighted below.

The Grafana user interface appears within the **Crosswork Manager** window, replacing the **Applications Detail** view.

Step 3 In the Grafana user interface, click **Home**. Grafana displays the list of monitoring dashboards and their categories, as shown in the following example.

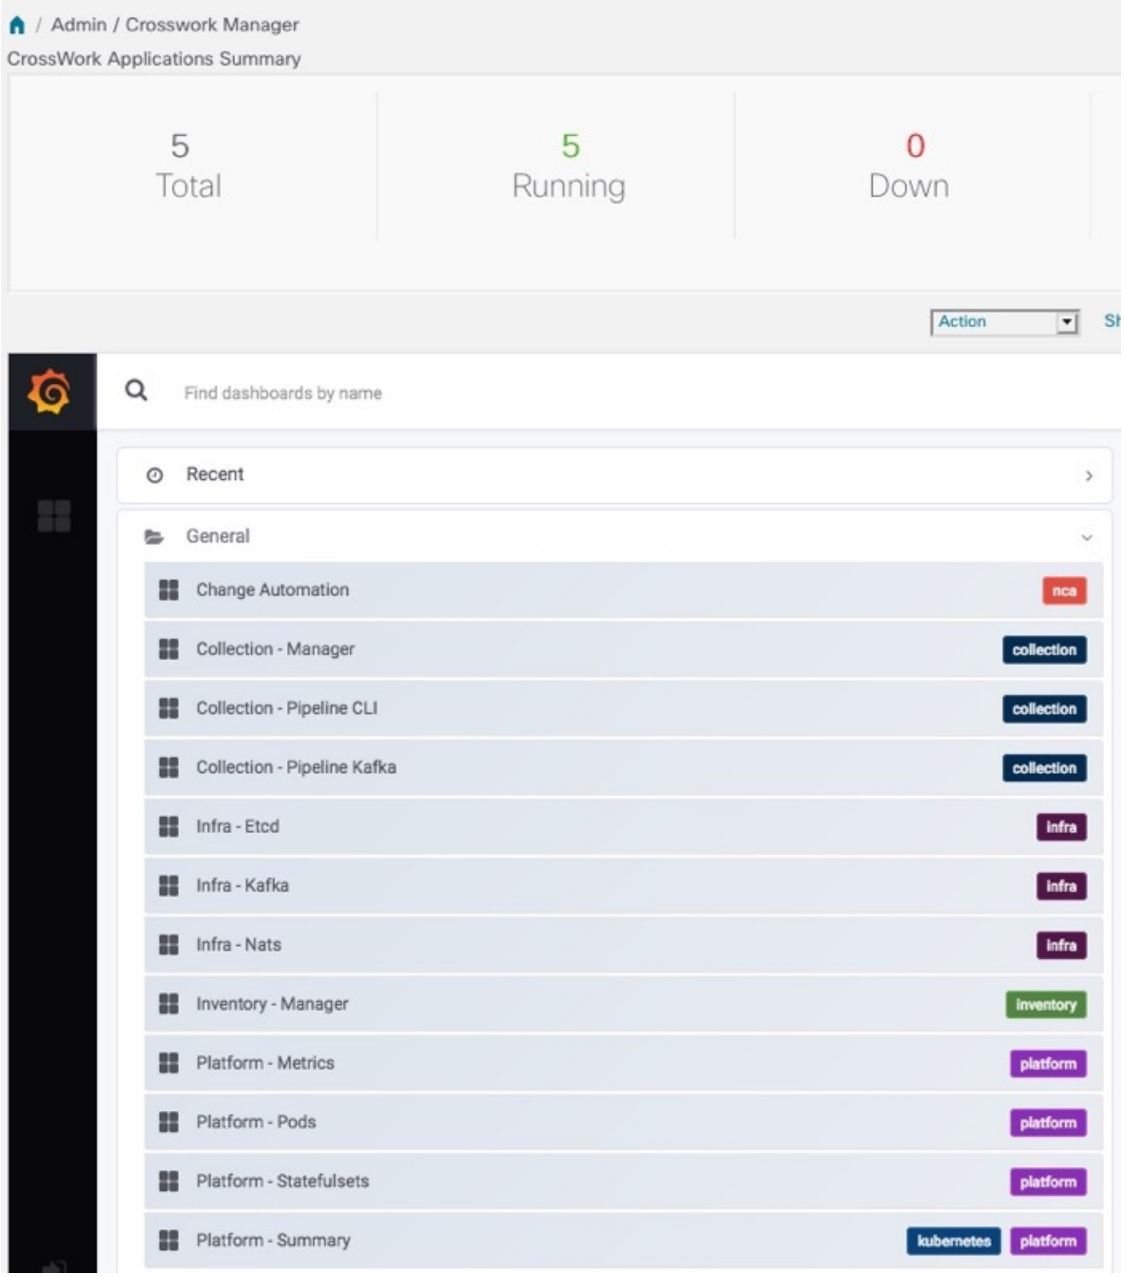

Step 4 Click the icon next to the dashboard you want to view. For example: Clicking on the **Platform - Summary** dashboard displays a view like the one shown in the following figure. For more information on how to use Grafana go to https://grafana.com.

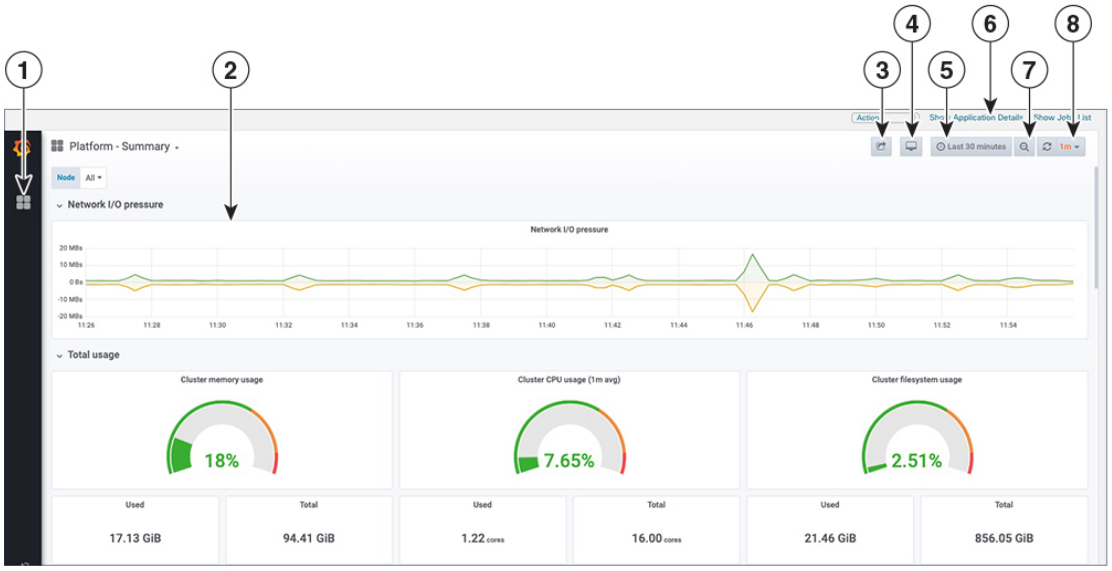

Step 5 Scroll the dashboard as needed to display all of the metrics it provides, or select any of the functions described in the following table.

Item	Description
1	**Dashboard Icon**: Click the icon to re-display the dashboard list and select a different dashboard.
2	**Time Series Graph Zoom**: You can zoom in on a specific time period within the graph of any time series data, as follows: **a.** Click a time-period starting point in the graph line and hold down the mouse. **b.** Drag the cursor to the endpoint. Light gray shading will appear in the block you are selecting. When you reach the endpoint, release the mouse. To reset a zoomed time series graph to the default, click the **Zoom Out icon**.
3	**Share Dashboard icon**: Click the icon to make the dashboard you are viewing shareable with other users. Clicking this icon displays a popup window with tabs and options to share the dashboard in your choice of these forms: • **URL Link**: Click the **Link** tab and then click **Copy** to copy the dashboard's URL to your clipboard. You can also choose whether to retain the current time and template settings with the URL. • **Local Snapshot File**: Click the **Snapshot** tab and then click **Local Snapshot**. Grafana creates a local snapshot of the dashboard on the server. When the snapshot is ready, click **Copy Link** to copy the URL of the snapshot to your clipboard. • **Export to JSON File**: Click the **Export** tab and then click **Save to file**. You will be prompted to save or open the exported JSON file. You can also choose to turn data source names in the file into templates by selecting the **Export for sharing externally** checkbox before clicking **Save to file**. • **View JSON File and Copy to Clipboard**: Click the **Export** tab and then click **View JSON** (you can choose to templatize data source names by selecting the **Export for sharing externally** checkbox before clicking **View JSON**). Grafana displays the exported JSON code in a popup window. Click **Copy to Clipboard** to copy the file to your clipboard.

Item	Description
4	**Cycle View Mode icon**: Click this icon to toggle between the default Grafana **TV** view mode and the **Kiosk** mode. The **Kiosk** view hides most of the Grafana menu. Press **Esc** to exit the **Kiosk** view.
5	**Time/Refresh Selector**: Indicates the time period for the metrics displayed in the dashboard and how often the metrics are refreshed. Click the selector to choose a different time range and refresh rate. You can specify a custom pair of time-range start and end points, or choose from one of several predefined ranges, such as **Today so far** or **Last three hours**. You can choose predefined refresh rates from **Off** to **2 Days**. When you have finished making changes, click **Apply**. When making selections, remember that Cisco Crosswork Optimization Engine keeps only 24 hours of data. If you select time ranges or refresh rates beyond that limit, the dashboard may be blank.
6	**Show Application Details**: Click this link to re-display the **Crosswork Manager** window's **Applications Detail** view.
7	**Zoom Out icon**: Click this icon to reset a zoomed time series graph back to the unzoomed state.
8	**Refresh icon**: Immediately refresh the data shown.

Collect and Share Cisco Crosswork Network Automation Logs and Metrics

You can collect logs and metrics on multiple levels of Cisco Crosswork Optimization Engine. You can collect logs and metrics for the entire system, for any of its installed application, or for any service supporting an application. You can also choose to collect only logs, only the additional metrics, or both.

Collected logs and metrics are stored in gzipped tar archive files. You can publish these archives to an HTTP or HTTPS server of your choice.

Step 1 From the main menu, choose **Admin** > **Crosswork Manager**. The **Crosswork Manager** window displays, with the **Application Detail** section listing all the applications.

Step 2 Click the option for the collection level and target information you want, as follows:

- To collect for the entire system: From the **Action** drop down on the right, opposite the **Applications Detail** section title, choose **Collect All**, **Collect Logs**, or **Collect Metrics**. See item 1 in the following figure.
- To collect for an application: Scroll to the **Application Detail** tile for the application you want. Then click the **Collect All**, **Collect Logs**, or **Collect Metrics** link on the right, opposite the application's name. See item 2 in the following figure.
- To collect for a service: Scroll to the **Application Detail** tile for the application whose service you want to collect. Click the **Show Service Detail** link for that application. Then click the **Collect All**, **Collect Logs**, or **Collect Metrics** link on the right, opposite the service's name. See item 3 in the following figure.

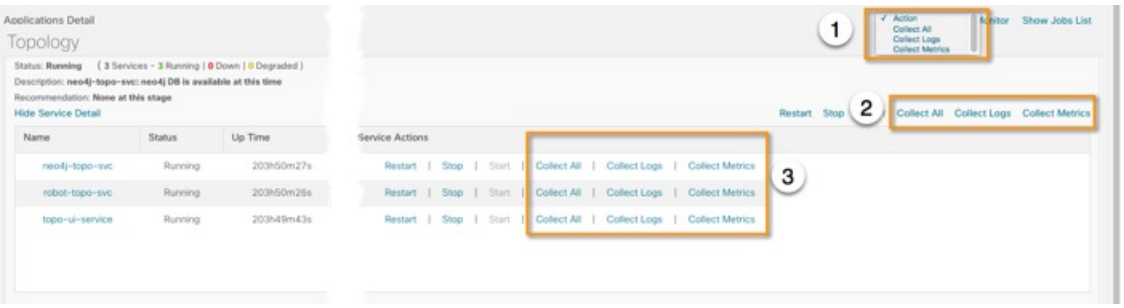

Step 3 When you click on the collection option you want, the **Crosswork Manager** window displays a popup message indicating that a job was successfully created and giving the job ID. Click on the **Show Jobs List** link at the right to view the job's progress in the **Crosswork Manager**window's **Jobs List** view, which replaces the **Applications Detail** view.

Step 4 Wait for the job to complete. When the **Jobs List** view's **Status** column for your job has changed to `JobCompleted`, the **Action** column for the job will show an enabled **Publish** link for the completed job, and the **Description** column will show the file name of the gzipped tar archive file containing the collected information.

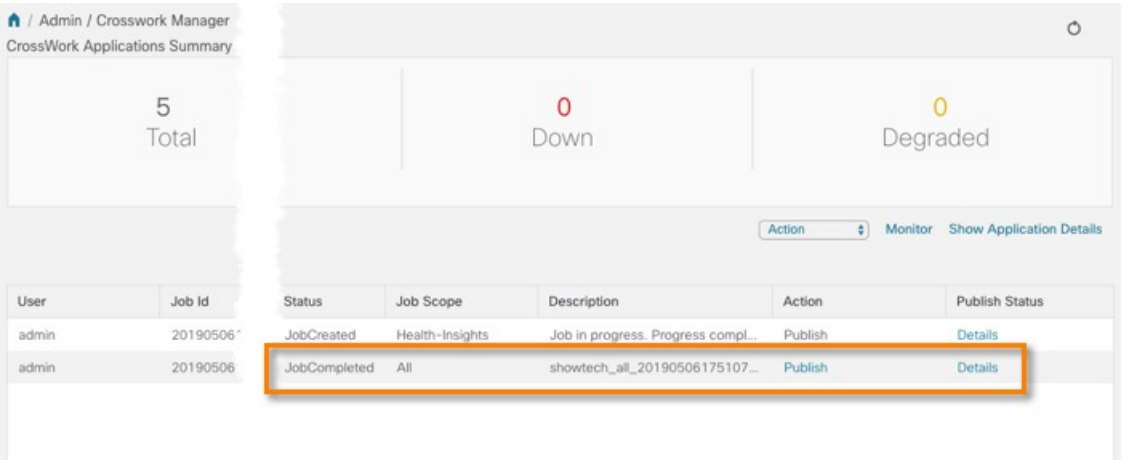

Step 5 (Optional) Click on the **Publish** link to publish the collected information to an HTTP or HTTPS server, as follows:

a) A popup window will prompt you for the destination server host name, the storage path on the server, the port number, and the login user name and password for the server (if required). Enter the server information and click **Publish**.

b) The **Job List** view's **Publish Status** column for the job shows an enabled **Details** link. Click the **Details** link to view a popup window showing the status of the publish job.

Step 6 When you are finished, click the **Show Application Details** link to re-display the **Applications Detail** view.

Control Cisco Crosswork Network Automation Applications and Services

Users with administrator privileges can control the runtime status of any Cisco Crosswork Optimization Engine application or service. This can include:

- Stopping a running application or service
- Starting a stopped application or service
- Restarting a running or stopped application or service

Please note that stopping, starting and restarting Cisco Crosswork Optimization Engine applications and services can result in anomalous system behavior and possible data loss. Use these functions only with the supervision of Cisco TAC staff.

Step 1 From the main menu, choose **Admin** > **Crosswork Manager**. The **Crosswork Manager** window displays, with the **Application Detail** view listing all the applications.

Step 2 Display the application or service whose runtime status you want to control:

- To control an application: Scroll to the **Application Detail** tile for the application you want.
- To control a service: Scroll to the **Application Detail** tile for the application whose service you want to control, then click the **Show Service Detail** link for that application to show its services.

Step 3 Click on the **Start**, **Stop**, or **Restart** link shown next to the service (item 1 in the following figure) or the application whose runtime status you want to control.

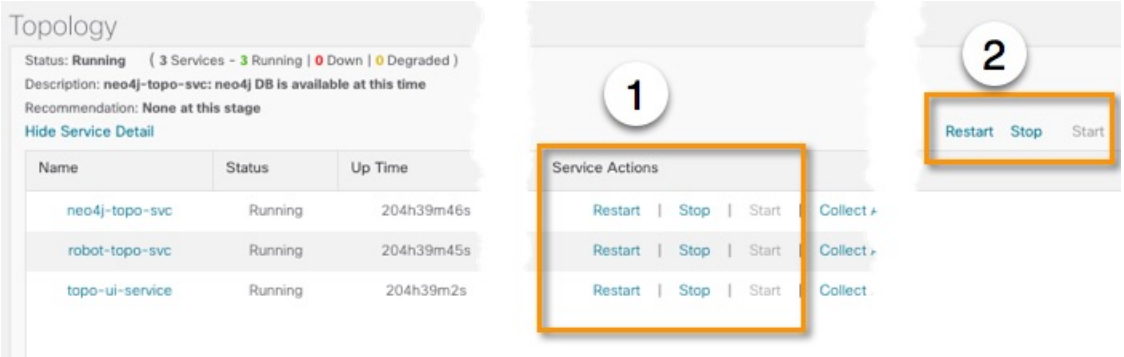

Step 4 Click the **Show Jobs List** link at upper right to view the runtime control job's progress in the **Crosswork Manager**window's **Jobs List** view.

Step 5 When you are finished, click the **Show Application Details** link to re-display the **Applications Detail** view.

Manage Backup and Restore

The Backup Restore functionality is critical to prevent data loss in your VM.

Follow the steps below to create a backup for the Cisco Crosswork Optimization Engine VM and to restore a backup.

Important

- Cisco recommends that you perform the backup or restore operation only during a scheduled maintenance window when admin users should not access the UI. Both operations are time-consuming and stops all other applications running in the system.
- The same Cisco Crosswork Optimization Engine software image that was used to backup must also be used when doing a restore operation.
- Stay on the **Backup Restore** window until the backup/restore process completes. Otherwise, you may see incorrect content or UI errors since various services are rebooting frequently.
- Only one backup or restore operation can be running at any given time.

Before you begin, ensure that:

- You have the Host Name, Port number, and Remote path to a Secure FTP server to use as the destination for backup files.
- You have the user credentials to an account with write permissions to create files and directories in the destination server remote path.

Step 1 From the main menu, choose **Admin** > **Backup Restore**. The **Backup Restore** window is displayed.

Step 2 During your first login, you should configure a destination server to store the backup file. This is a one-time activity and has to be completed before taking the backup. Click **Destination** to display the **Edit Destination** dialog box. Make relevant entries in the fields provided.

Click **Save** to confirm the server details.

Step 3 **To create a backup:**

a) Click **Backup**. The **Backup** dialog box is displayed with destination server details pre-filled.
b) Provide a relevant name in the **Job Name** field.
c) (Optional) Click **Verify Backup** to check if Cisco Crosswork Optimization Engine has enough resources to complete the operation. If the check is successful, a warning message is displayed about the time-consuming nature of the operation. Click **OK**.
d) Click **Start Backup** to start the backup operation. The corresponding backup job set is created and added to the job list. See Step 5 to view Backup progress.

Step 4 **To restore a backup file:**

a) Select the required backup file from the **Backup Restore Job Sets** table, and the job details are displayed on the right side.
b) Click the **Restore** button to display the **Restore** dialog box with destination server details pre-filled.
c) Provide a relevant name in the **Job Name** field.
d) (Optional) Click **Verify Restore** and a prompt is displayed that suggests doing the backup or restore during maintenance window owing to the time-consuming nature of the operation. Click **OK**.
e) Click **Start Restore** to start the restore operation. The corresponding restore job set is created and added to the job list.

Step 5 **To view a job progress:**

a) Enter the job details (such as Status, Job Name, or Job Type) in the search fields in **Backup Restore Job Sets** table on the left side. Click to select which columns to display in the Job set list. The list is automatically filtered based on your search string. Click the required job set from the search results.
b) Alternately, you can manually scroll the list and click the required job set.
c) The **Job Details** table on the right side displays information about the selected job set such as Status, Job Type and Start time. In case of a failed job, hover the mouse pointer over the icon near **Status** to view the error details.

Disaster Restore

Disaster Restore is a restore operation, appropriately named to be used in case of a disaster, such as VM crash. The **Disaster Restore** option is displayed if no backup jobs have been initiated in the system. After the completion of the first backup job, this button is disabled.

Note While using disaster recovery operation, please note the following:

- The new VM that you use needs to have the same IP address as the one where backup was performed. This is important as internal certificates are tied to the IP address.
- The same Cisco Crosswork Optimization Engine software image that was used to backup must also be used when doing a restore operation.
- The VM which is brought up should have same services running when the backup was performed. If the previous VM was patched/upgraded then the new VM also needs to be patched/upgraded before disaster restore is performed.
- The disaster restore operation trusts the backup file which is provided. Caution is advised while selecting the appropriate backup file.

To perform a disaster restore:

Step 1 From the main menu, choose **Admin** > **Backup Restore**. The **Backup Restore** window is displayed.

Step 2 Click **Destination** to display the **Edit Destination** dialog box. Enter the details of the remote destination server where the backup file is uploaded.

Step 3 Click **Disaster Restore** to display the **Disaster Restore** dialog box with destination server detailed pre-filled.

Step 4 Make relevant entry in the **Backup File Name** field.

Step 5 Click **Start Restore** to start the disaster restore operation.

Note
- If disaster restore operation fails, you are recommended to bring up a new VM to retry the disaster restore operation.
- If you find that there are missing SR policies or RSVP-TE tunnels in your topology, use the Configuration Database CLI tool (see Configuration Database CLI Tool, on page 137).

Configuration Database CLI Tool

The Configuration Database contains all SR policies and RSVP-TE tunnels that Cisco Crosswork Optimization Engine is aware of. The Configuration Database is updated whenever an SR policy or RSVP-TE tunnel is provisioned, modified, or deleted. The Configuration Database CLI tool is a utility that can do the following:

- Reads/writes CSV files to the Configuration Database
- Populates SR policy and RSVP-TE tunnel information from the Configuration Database to create a CSV file

The Configuration Database CLI tool is especially useful when trying to recover missing SR policies and RSVP-TE tunnels after a Restore operation. For example, the `--dump-missing` option produces a CSV file which lists missing SR policies and RSVP-TE tunnels. After reviewing this CSV file, you determine which SR policies and RSVP-TE tunnels were provisioned by Cisco Crosswork Optimization Engine and load them back into the topology using the `--load` option. See the CLI tool help for more information.

Step 1 Enter the **optima-pce-dispatcher** container:

```
kubectl exec -it optima-pce-dispatcher-XXXXXXX-XXXX bash
```

Step 2 You can run the following commands:

a) Show CLI tool help text.

```
python3 /opt/optima/pce_dispatcher/config_db/csv_util.py --help
```

b) Save all SR policies and RSVP-TE tunnels that are in the Configuration Database to a CSV file.

```
python3 /opt/optima/pce_dispatcher/config_db/csv_util.py --dump=/path/to/file/dump_file.csv
```

c) Load the contents from the provided CSV file and write policies to the Configuration Database.

```
python3 /opt/optima/pce_dispatcher/config_db/csv_util.py --load=/path/to/file/load_file.csv
```

Note If any duplicate SR policies (combination of headend, endpoint, and color) or RSVP-TE tunnels (combination of headend and tunnel name) are found, they will be overwritten. Only valid TE tunnels will be added to the Configuration Database.

d) Compare SR policies and RSVP-TE tunnels that are currently in the topology with what is saved in the Configuration Database and save the missing SR policies and RSVP-TE tunnels to a CSV file. This CSV file can then be used to load the missing policies into the Configuration Database.

```
python3 /opt/optima/pce_dispatcher/config_db/csv_util.py -dump-missing=/path/to/file/dump_file.cs
```

Manage Users

From the main menu, select **Admin** > **Users** to display the **Users** window. Using this window, you can add a new user, edit the settings for an existing user, delete a user from the network, and create user roles.

Note Before you can create a new user that does *not* have admin-level access to Cisco Crosswork Optimization Engine functionality, you must first create a new role that limits the features they can access. See Create User Roles for more information.

Only a local admin user can add, update, and delete other local user accounts. A TACACS+ user, regardless of role assigned, will not be able to manage local users.

Administrative Users Created During Installation

During installation, Cisco Crosswork Optimization Engine creates two special administrative IDs:

1. The **virtual machine administrator**, with the username **`cw-admin`**, and the default password**`cw-admin`** . Data center administrators use this ID to log in to and troubleshoot the VM hosting the Cisco Crosswork Optimization Engine server.
2. The **Crosswork administrator**, with the username **`admin`** and the default password **`admin`**. Product administrators use this ID to log in to and configure the Cisco Crosswork Optimization Engine user interface, and to perform special operations, such as creating new user IDs.

The default password for both administrative user IDs must be changed the first time they are used. You can also change the Crosswork administrator password using the following methods:

- Log in as the admin user and edit the admin user password, as explained in Edit Users, on page 139.
- Enter the following command: `admin(config)# **username admin** <password>`

Add Users

Follow the steps below to create a new user ID.

The user ID's user name must be unique. You cannot create a new user ID with the same user name as an existing user ID.

The special administrative user names **`admin`** (for administering Cisco Crosswork Optimization Engine) and **`cw-admin`** (for administering the virtual machine hosting the product) are created during installation and are reserved for those purposes (see Administrative Users Created During Installation, on page 138).

Step 1 From the main menu, choose **Admin** > **Users**.

The **Users** window opens.

If it is not already displayed, click the **Users** tab.

Step 2 Click [+] to open the **Add New User** dialog box.

Step 3 Enter the following information for the user you are adding:

- **User Name**: Enter a unique name for the user ID. User names cannot contain spaces or special characters.
- **First Name** and **Last Name**: Enter the first and last name of the person assigned to this user ID.

- From the **Role** drop-down at the bottom of the dialog box, choose the role that you want to assign to the user. See Create User Roles for more information.
- **Password** and **Confirm Password**: Enter the default password for this user ID. The user will be required to change the default password the first time they attempt to log on using it.

Note The user password must be string of minimum 8 characters without spaces and should include letters, numbers, upper-case and lower-case characters, and one of the allowed special characters ("@!$%*?&").

Step 4 Click **Save**.

Edit Users

Users with administrator privileges can edit any user ID's User Name, First Name, Last Name, and Role.

Administrators cannot change a user's password by editing the user ID. Users can change their passwords by logging in, clicking , and selecting **Change Password**.

Step 1 From the main menu, choose **Admin** > **Users**.

The **Users** window opens.

If it is not already displayed, click the **Users** tab.

Step 2 Click on the check box of the user whose settings you want to update, then click to open the **Edit User** dialog box.

Step 3 Make the necessary updates to the user ID.

Note First Name, Last Name and Role can be edited for user accounts with administrative privileges.

Step 4 Click **Update** to save your changes.

Delete Users

Follow the steps below to delete an existing user ID.

The administrative user IDs **`admin`** and **`cw-admin`** created during installation cannot be deleted (see Administrative Users Created During Installation, on page 138).

Step 1 From the main menu, choose **Admin** > **Users**.

The **Users** window opens.

If it is not already displayed, click the **Users** tab.

Step 2 Click on the check box of the user you want to delete, then click . The**Delete** *Username* **User** dialog displays.

Step 3 Click **Delete** to confirm deletion.

Create User Roles

Local users with administrator privileges can create new users as needed (see Add Users, on page 138).

Users created in this way can perform only the functions or tasks that are associated with the user role they are assigned.

The local **admin** role enables access to all functionality. It is created during installation and cannot be changed or deleted. However, its privileges can be assigned to new local users. Only local users can create or update user roles; TACACS users cannot.

Follow the steps below to create a new user role.

Step 1 From the main menu, choose **Admin** > **Users**.

The **Users** window opens.

If it is not already displayed, click the **Roles** tab. The **Roles** window has a **Roles** table on the left side and a corresponding **admin** table on the right side which shows the grouping of user permissions for the selected role.

Step 2 On the **Roles** table, click [+] to display a new role entry in the table.

Step 3 Enter a unique name for the new role.

Step 4 Define the user role's privilege settings:

a) Check the check box for every API that users with this role can access. The APIs are grouped logically based their corresponding application.

b) For each API, define whether the user role has **Read**, **Write**, and **Delete** permission by checking the appropriate check box. You can also select an entire API group (such as AAA), and all the APIs under the group will be selected with **Read**,**Write** and **Delete** permissions pre-selected.

Step 5 Click **Save** to create the new role.

To assign the new user role to one or more user IDs, edit the **Role** setting for the user IDs (see Edit Users, on page 139).

Edit User Roles

Users with administrator privileges can quickly change the privileges of any user role other than the default **admin** role.

Step 1 From the main menu, choose **Admin** > **Users**.

The **Users** window opens.

If it is not already displayed, click the **Roles** tab.

Step 2 In the **Roles** table, click on an existing role to select it. The **Admin** table on the right side displays the permission settings for the selected role.

Step 3 Define the role's settings:

a) Check the check box for every API that the role can access.

b) For each API, define whether the role has **Read**, **Write**, and **Delete** permission by checking the appropriate check box. You can also select an entire API group (such as AAA), and all the APIs under the group will be selected with **Read**,**Write** and **Delete** permissions pre-selected.

Step 4 When you are finished, click **Save** to save your changes.

Clone User Roles

Cloning an existing user role is the same as creating a new user role (see Create User Roles, on page 140), except that you need not set privileges for it. If you like, you can let the cloned user role inherit all the privileges of the original user role.

Cloning user roles is a handy way to create and assign many new user roles quickly. Following the steps below, you can clone an existing role multiple times. Defining the cloned user role's privileges is an optional step; you are only required to give the cloned role a new name. If you like, you can assign it a name that indicates the role you want a group of users to perform. You can then edit the user IDs of that group of users to assign them their new role (see Edit Users, on page 139). Later, you can edit the roles themselves to give users the privileges you want (see Edit User Roles).

Step 1 From the main menu, choose **Admin** > **Users**.

The **Users** window opens.

If it is not already displayed, click the **Roles** tab.

Step 2 Click on an existing role to select it.

Step 3 Click to create a new duplicate entry in the **Roles** table with all the permissions of the original role.

Step 4 Enter a unique name for the cloned role.

Step 5 (Optional) Define the role's settings:

a) Check the check box for every API that the cloned role can access.

b) For each API, define whether the clone role has **Read**, **Write**, and **Delete** permission by checking the appropriate check box. You can also select an entire API group (such as AAA), and all the APIs under the group will be selected with **Read**,**Write** and **Delete** permissions pre-selected.

Step 6 Click **Save** to create the newly cloned role.

Delete User Roles

Users with administrator privileges can delete any user role that is not the default **admin** user role or that is not currently assigned to a user ID. If you want to delete a role that is currently assigned to one or more user IDs, you must first edit those user IDs to assign them to a different user role.

Step 1 From the main menu, choose **Admin** > **Users**.

The **Users** window opens.

If it is not already displayed, click the **Roles** tab.

Step 2 Click on the role you want to delete, to select it.

Step 3 Click to display the **Delete Role** dialog box.

Step 4 Click **Delete** to confirm that you want to delete the user role.

Manage TACACS+ Servers

In addition to local database authentication, Cisco Crosswork Optimization Engine can use TACACS+ servers to authenticate users. TACACS+ is a security protocol that provides centralized validation of users attempting to access your network. It allows for a single access control server (the TACACS+ daemon) to provide authentication, authorization, and accounting (AAA) services independently of one another.

Local database authorization takes precedence over authorization by TACACS+ server. When adding the TACACS+ server, you can specify the priority value for each instance.

Note Please note that any operation you do following the instructions in this section will affect all new logins to the Cisco Crosswork Optimization Engine user interface. To minimize session interruption, Cisco recommends that you perform all your TACACS+ changes and submit them in a single session.

Add a TACACS+ Server

Before adding a TACACS+ server, you will need to know the server's IP address, port number, shared secret, and service name.

Step 1 From the main menu, choose **Admin** > **AAA**.

The **AAA** window opens. If it is not already displayed, click the **TACACS+ Servers** tab.

Step 2 Click to open the **Add Server** dialog box.

Step 3 Enter the TACACS+ server's settings, then click **Add**.

Note Only the server's IP address, port number, shared secret, and service name are required. You can leave the other values blank, as needed.

Step 4 Click **Save Server Changes** to submit the changes. You will be prompted with a warning message about restarting the server to update the changes. Click **Save Changes** to confirm.

Edit a TACACS+ Server

Step 1 From the main menu, choose **Admin** > **AAA**.

The **AAA** window opens. If it is not already displayed, click the **TACACS+ Servers** tab.

Step 2 Click the check box next to the TACACS+ server whose settings you want to update, then click .

The **Edit Server** dialog box opens.

Step 3 Make the necessary changes, then click **Update**.

Note You cannot change the value for the **Shared Secret** parameter.

Step 4 Click **Save Server Changes** to submit the changes. You will be prompted with a warning message about restarting the server to update the changes. Click **Save Changes** to confirm.

Delete a TACACS+ Server

Step 1 From the main menu, choose **Admin** > **AAA**.

The **AAA** window opens. If it is not already displayed, click the **TACACS+ Servers** tab.

Step 2 Click the check box next to the TACACS+ server you want to delete.

Note You can delete only one TACACS+ server at a time.

Step 3 Click . The **Delete** *server-IP-address* dialog box opens.

Step 4 Click **Delete** to confirm.

Manage LDAP Servers

Cisco Crosswork Optimization Engine supports the use of LDAPv3 servers with OpenLDAP to authenticate users. Lightweight Directory Access Protocol (LDAP) is a server protocol used to access and manage directory information. It manages directories over IP networks and runs directly over TCP/IP using simple string formats for data transfer.

Like TACACS+ server, you can specify the priority value to assign precedence in the authentication request.

Note Please note that any operation you do following the instructions in this section will affect all new logins to the Cisco Crosswork Optimization Engine user interface. To minimize session interruption, Cisco recommends that you perform all your LDAP changes and submit them in a single session.

Add a LDAP Server

Before adding a LDAP server, you will need to know the Server name and URL, Bind DN and credential, Base DN, user filter, DN format, Principal Attribute ID, Policy ID, and connection timeout value.

Before you begin

Note the following:

- Cisco Crosswork Optimization Engine supports the use of LDAPv3 servers with OpenLDAP to authenticate users
- Roles needs to be mapped exactly to the same LDAP CrossworkPolicyId. See the example figure below.
- The user name in Crosswork and LDAP cannot be the same. If they are, the Crosswork user is prioritized.
- LDAP users cannot create Roles/Users even if it has an admin role.
- There are no errors that will indicate a misconfiguration when adding an LDAP server.

Step 1 From the main menu, choose **Admin** > **AAA**.

The **AAA** window opens. Click on the **LDAP Servers** tab.

Step 2 Click ⊞ to open the **Add Server** dialog box.

Step 3 Enter the LDAP server settings, then click **Add**.

Step 4 Click **Save Server Changes** to submit the changes. You will be prompted with a warning message about restarting the server to update the changes. Click **Save Changes** to confirm.

The following figure shows a sample OpenLDAP configuration:

Figure 54: Adding an LDAP Server (OpenLDAP Configuration)

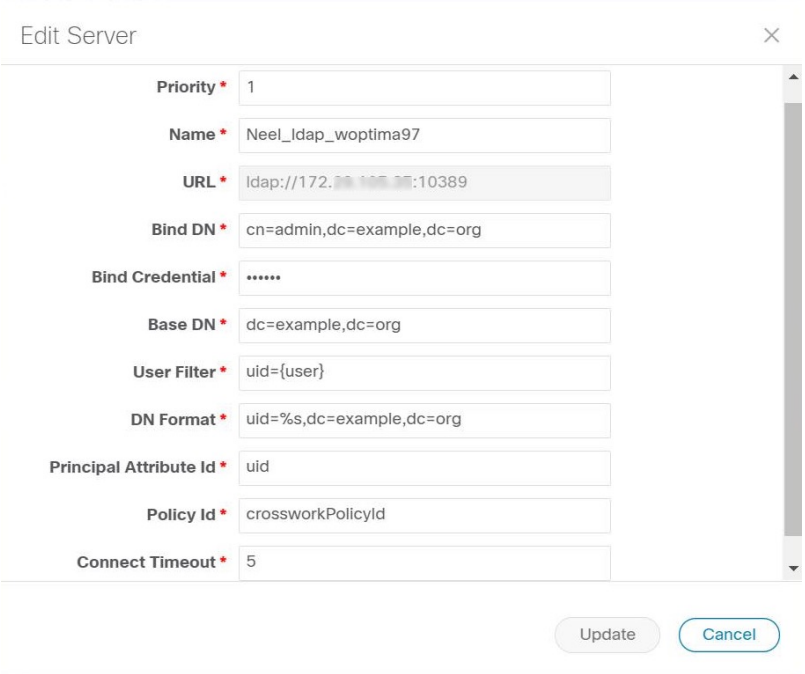

Edit a LDAP Server

Step 1 From the main menu, choose **Admin** > **AAA**.

The **AAA** window opens. Click on the **LDAP Servers** tab.

Step 2 Click the check box next to the LDAP server whose settings you want to update, then click .

The **Edit Server** dialog box opens.

Step 3 Make the necessary changes, then click **Update**.

Step 4 Click **Save Server Changes** to submit the changes. You will be prompted with a warning message about restarting the server to update the changes. Click **Save Changes** to confirm.

Delete a LDAP Server

Step 1 From the main menu, choose **Admin** > **AAA**.

The **AAA** window opens. Click on the **LDAP Servers** tab.

Step 2 Click the check box next to the LDAP server you want to delete.

Note You can delete only one LDAP server at a time.

Step 3 Click . The **Delete** *server-IP-address* dialog box opens.

Step 4 Click **Delete** to confirm.

Define Network Topology Display Settings

Cisco Crosswork Optimization Engine administrator privileges are required to configure the display settings that are used by the Network Topology application.

For a description of how to configure these settings, see the following topics:

- Define Color Thresholds for Link Bandwidth Utilization
- Configure Geographical Map Settings, on page 69

Manage Certificates

The Cisco Crosswork Optimization Engine VM-hosted server and its browser-based user interface communicate with each other using SSL certificates exchanged over HTTPS. For details about these protocols, see SSL Certificates, on page 154 and HTTPS, on page 154

When installed, Cisco Crosswork Optimization Engine secures these interactions using a self-signed TLS certificate. This certificate has a two-year lifespan, after which it expires. If you want to continue using the expired self-signed certificate to secure server/client communications, you will need to regenerate it by following the steps in Extend Self-Signed Certificate Expiration, on page 147

If you prefer to secure these communications with a user-provided certificate, either purchased from a Certificate Authority (CA) or self-signed by your organization, you can validate and upload it by following the steps in Substitute a User-Provided Certificate, on page 147.

The user-provided certificate must meet the following requirements:

- Cisco Crosswork Optimization Engine supports IP Subject Alternative Name (SAN) server certificates only. The IP address is the primary means to reach the user interface.
- The server will present your user-provided certificates to the browser, so the certificates you supply must be valid both for Cisco and for Cisco Crosswork Optimization Engine.
- It must also include the required fields and field values shown in the following table.

Table 15: Required User-Provided Certificate Fields and Values

Field	Description	Value
<NUMBER OF DAYS>	Number of days the certificate will be valid.	Must be greater than **30** days and less than **730** days (or two years)
<COUNTRY>	Country (`C=`)	**US**
<STATE>	State (`ST=`)	**CALIFORNIA**
<LOCATION>	Location (`L=`)	**SAN JOSE**
<ORGANIZATION>	Organization (`O=`)	**CISCO SYSTEMS INC**
<ORGANIZATIONAL UNIT NAME>	Organizational Unit (`OU=`)	**CROSSWORK**
<COMMON NAME>	Common Name (`CN=`)	The IP address of the Cisco Crosswork Optimization Engine server VM.

- The certificate must also have the SAN extension set, with both DNS and IP address keys. The following provides an example of how to generate a self-signed certificate using OpenSSL:

```
/usr/bin/openssl req \
                   -x509 \
                   -nodes \
                   -days 730 \
                   -newkey rsa:4096 \
                   -keyout "filename.key" \
                   -out "filename.crt" \
                   -subj "/C=US/ST=CALIFORNIA/L=SAN JOSE/O=CISCO SYSTEMS
INC/OU=CROSSWORK/CN=1.1.1.1" \
                   -extensions SAN \
                   -config <(cat /etc/ssl/openssl.cnf \
                    <(printf "\n[SAN]\nsubjectAltName=DNS:0.0.0.0,IP:1.1.1.1"))
```

Extend Self-Signed Certificate Expiration

Follow these steps to regenerate the self-signed certificate and extend its lifetime by two years.

Step 1 From the main menu, select **Admin** > **Certificate Management**. The **Certificate Management** window appears.

Step 2 Select the **Regenerate Crosswork Certificate** radio button.

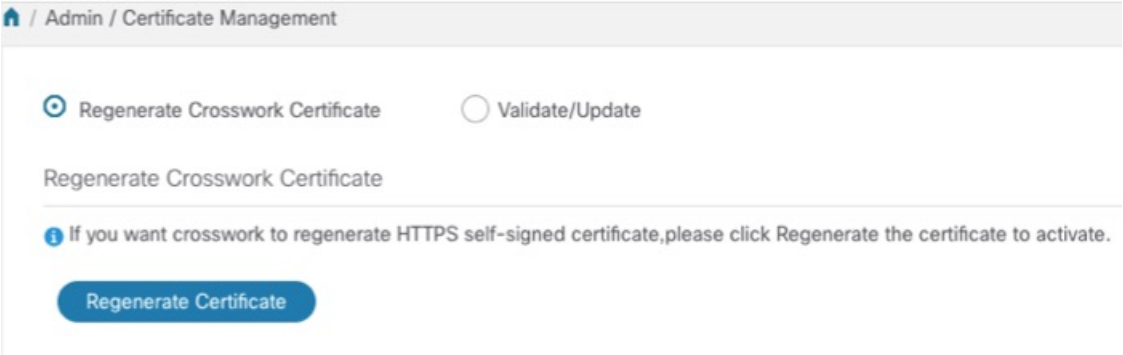

Step 3 When you are ready, click **Regenerate Certificate**.

When Cisco Crosswork Optimization Engine has finished regenerating the certificate, it displays an alert message indicating that the regeneration operation is successful and you will be logged out. You must log in again to continue using Cisco Crosswork Optimization Engine.

Substitute a User-Provided Certificate

Follow the steps below to validate and upload a user-provided certificate. The certificate must meet the requirements explained in Manage Certificates, on page 145.

Before you begin

You must know the names of the user-provided certificate and key files and their locations in your local storage.

Step 1 From the main menu, select **Admin** > **Certificate Management**. The **Certificate Management** window appears.

Step 2 Select the **Validate/Update** radio button.

Step 3 Use the **Browse** button next to each field to browse to and select the key and certificate files you want to validate and use.

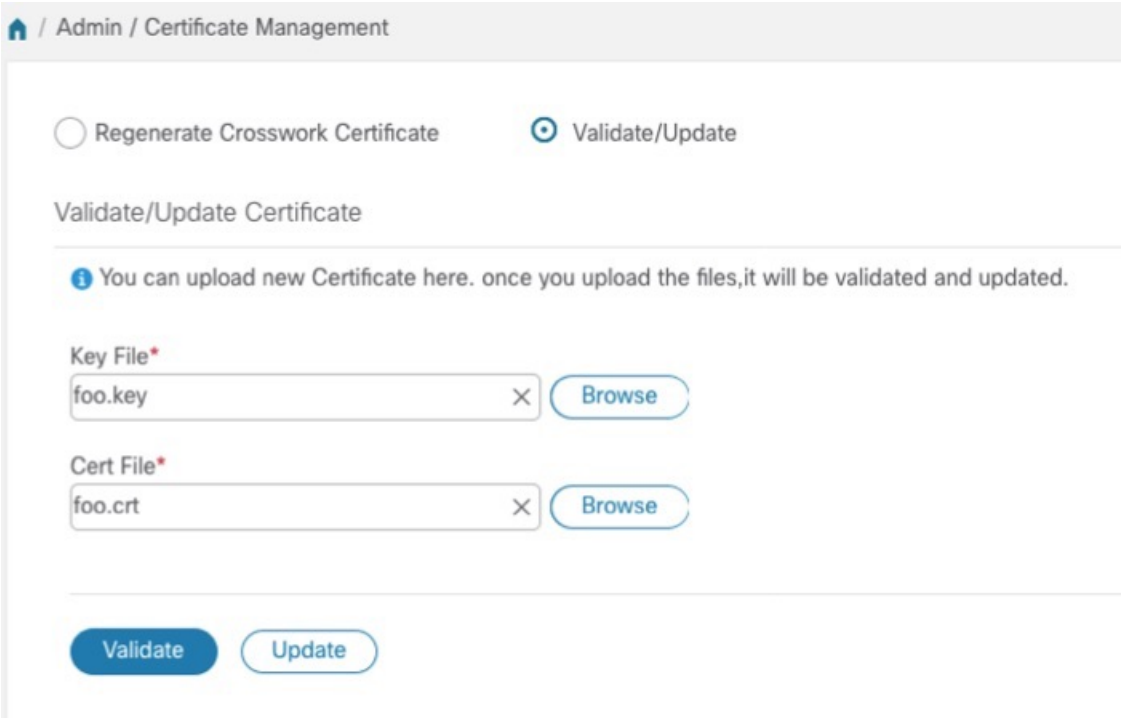

Step 4 Click **Validate** to validate the certificate and key files.

Step 5 Click **Update** to replace the existing certificate with the user-provided certificate you have validated.

Smart Licensing Registration

This section provides an overview of the Cisco Smart Licensing feature integrated with the and describes the instructions to complete the product registration.

Overview

Smart Licensing is a software based end-to-end license platform that comprises several tools and processes that authorizes customers to use Cisco products. Smart Licensing provides a software inventory management system that provides Customers, Cisco, and selected Partners with information about Software Ownership and Software Utilization.

A **Cisco Smart Account** provides the repository for Smart enabled products and enables you to activate Cisco licenses, monitor license usage and track Cisco purchases. The **Cisco Smart Software Manager (CSSM)** enables you to manage all your Cisco Smart software licenses from one centralized website. With Cisco Smart Software Manager, you may create and manage multiple virtual accounts within your Smart Account to manage licenses. For more information, see https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager.html

From the main menu, select **Admin** > **Smart Licensing Registration** to display the **Smart Software Licensing** window. Using this window, you can register your application, edit the transport settings, renew the license, and de-register your application.

Prerequisites for Smart Licensing Registration

You should have:

- A Cisco Smart Account.
- Purchased licenses for the Cisco Crosswork Optimization Engine application.

Configure Transport Settings

You can configure the transport settings to decide how communicates with the Cisco servers.

- **Direct**: The application directly connects with Cisco Smart Software Manager (CSSM).
- **Transport Gateway**: The application communicates via a Transport Gateway or CSSM on-prem, which replicates the cloud-based user experience but keeps all communication on premises.

  Note For more information on the CSSM on-prem option, see the Smart Software Manager guide.

- **HTTP/HTTPS Gateway**: The application connects via an intermediate proxy server. This is applicable only for Direct mode.

Note Transport Settings cannot be changed while the is in Registered mode. You have to de-register to change them.

Step 1 In the **Smart Software Licensing** window, the Transport Settings display the current transport mode selected. To modify, click **View/Edit**.

The **Transport Settings** dialog box is displayed.

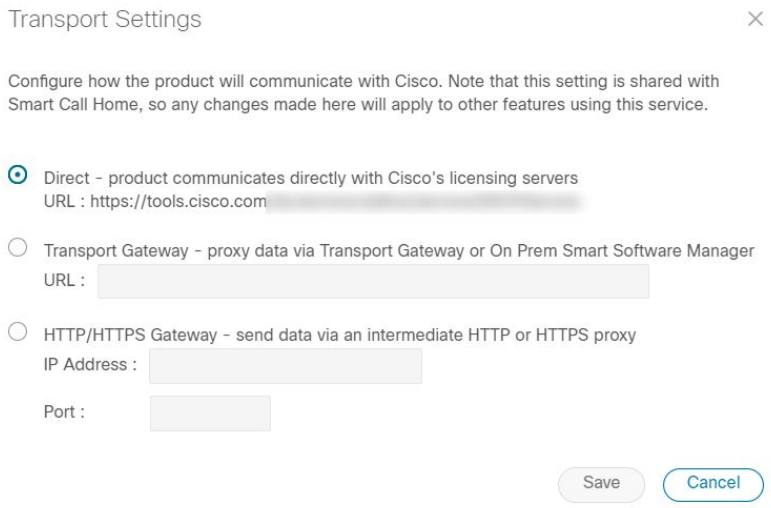

Step 2 Select the relevant transport mode and make relevant entries in the fields provided.

Step 3 Click **Save**.

Register

To enable licensed features, must be registered to CSSM using a registration ID token. Once registered, an Identity Certificate is saved securely in the Smart Account and used for all ongoing communications. The certificate is valid for one year and will be renewed automatically after six months to ensure continuous operation.

Note For information on generating the registration token, please refer to the support resources provided in the Smart Software Manager webpage.

Step 1 From the main menu, select **Admin** > **Smart Licensing Registration** to display the **Smart Software Licensing** window. The registration status

The registration status and license authorization status will be **Unregistered** and **Evaluation mode** respectively.

Figure 55: Smart Software Licensing Unregistered

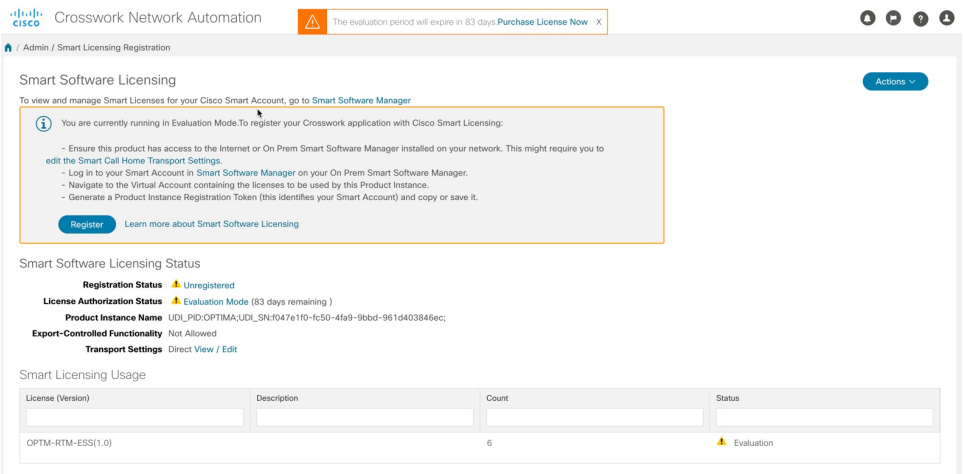

Step 2 In the **Smart Software Licensing** window, click **Register**.

The **Smart Software Licensing Product Registration** dialog box is displayed.

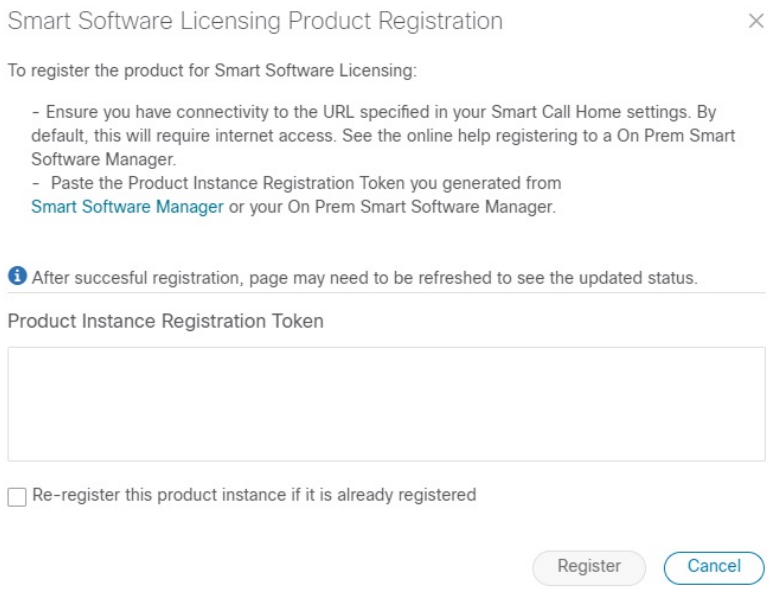

Step 3 In the **Product Instance Registration Token** field, enter the registration token generated from your Smart Account. Make sure the token ID is accurate and within validity period. For more information, see https://www.cisco.com/c/en_in/products/software/smart-accounts/software-licensing.html.

Step 4 (Optional) If you are re-registering the application, check the **Re-register this product registration if it is already registered** checkbox.

Note After a backup restore or disaster restore operation, you must manually re-register the VM to CSSM. This is applicable in case of a VM that has been already registered while taking the backup which is used in the restore operations.

Step 5 Click **Register**. It may take a few minutes to process the registration. If successful, the 'Product Registration completed successfully' message is displayed.

The registration status and license authorization status will be updated as **Registered** and **Authorized** respectively.

Note
- If you encounter a communication timeout error during registration, click **OK** in the error dialog box and the application will reattempt the registration.
- In some cases, after succesful registration, the page may need to be refreshed manually to see the updated status.

Manual Actions

The renewal of registration and authorization are automatically enabled for , by default. However, in the event of a communication failure between the application and the Cisco server, these actions can be manually initiated. You can use the **Actions** drop-down button to manually renew, re-register and de-register the application.

Step 1 In the **Smart Software Licensing** window, click **Actions** drop-down button and select the relevant option for the following quick actions.

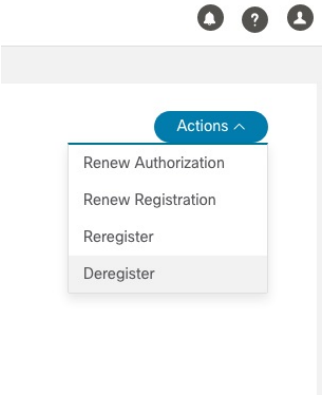

a) **Actions** > **Renew Authorization**: To renew the authorization manually if the automatic renewal service fails at the end of 30 days.
b) **Actions** > **Renew Registration**: To renew the registration manually if the automatic renewal service fails at the end of 6 months.
c) **Actions** > **Re-register**: Re-register the application, for example, on account of the expiry of registration tokens.
d) **Actions** > **De-register**: De-register the application, for example, when the transport settings need to be changed.

Note Once de-registered, the application will be moved to **Evaluation** mode (if evaluation period is available), or **Evaluation Expired** mode. For more information, see License Authorization Statuses, on page 153

Step 2 The selected action is executed successfully.

License Authorization Statuses

Based on the registration status of your application, you can see the following License Authorization Statuses.

Table 16: License Authorization Statuses

Registration Status	License Authorization Status	Description
Unregistered	Evaluation mode	A 90-day evaluation period during which the licensed features of the application can be freely used. This state is initiated when you use the application for the first time.
	Evaluation Expired	The application has not been successfully registered at the end of the evaluation period. During this state, the application features are disabled, and you must register to continue using the application.
	Registered Expired	The application is unable to contact the CSSM before the expiration of Identity Certificates and has returned to the unregistered state. The application resumes the remaining evaluation period, if available. At this stage, new registration ID token is required to reregister the application.
Registered	Authorized (In Compliance)	The application has been fully authorized to use the reserved licensed features. The authorization is automatically renewed every 30 days.
	Out of Compliance	The associated Virtual Account does not have enough licenses to reserve for the application's current feature use. You must renew the entitlement/usage limit registered with the token to continue using the application. See figure below.
	Authorization Expired	The application is unable to communicate with the CSSM for 90 days or more, and the authorization has expired.

Figure 56: Registered and Out of Compliance Sample

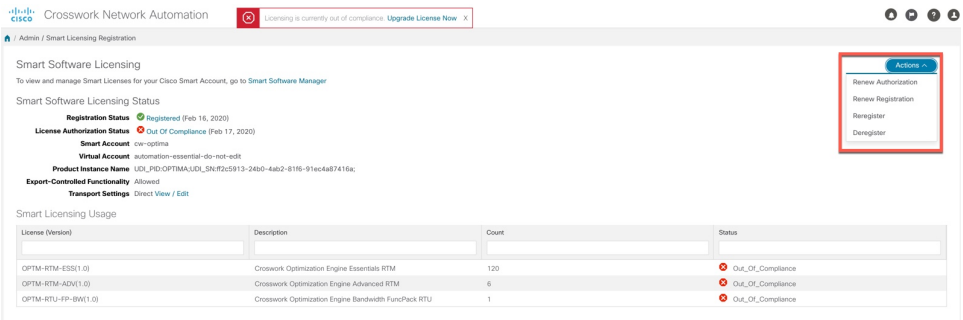

Security Hardening Overview

Security hardening entails making adjustments to ensure that the following components optimize their security mechanisms:

- infrastructure
- storage system (local or external)

Hardening security requires completion of the following tasks:

- Shutting down insecure and unused ports
- Configuring network firewalls
- Hardening the infrastructure, as needed

Although your primary source of information is your Cisco representative, who can provide server hardening guidance specific to your deployment, you can also follow the steps in this section to secure .

Core Security Concepts

If you are an administrator and are looking to optimize the security of your product, you should have a good understanding of the following security concepts.

HTTPS

Hypertext Transfer Protocol Secure (HTTPS) uses Secure Sockets Layer (SSL) or its subsequent standardization, Transport Layer Security (TLS), to encrypt the data transmitted over a channel. Several vulnerabilities have been found in SSL, so now supports TLS only.

Note TLS is loosely referred to as SSL often, so we will also follow this convention.

SSL employs a mix of privacy, authentication, and data integrity to secure the transmission of data between a client and a server. To enable these security mechanisms, SSL relies upon certificates, private-public key exchange pairs, and Diffie-Hellman key agreement parameters.

SSL Certificates

SSL certificates and private-public key pairs are a form of digital identification for user authentication and the verification of a communication partner’s identity. Certificate Authorities (CAs), such as VeriSign and Thawte, issue certificates to identify an entity (either a server or a client). A client or server certificate includes the name of the issuing authority and digital signature, the serial number, the name of the client or server that the certificate was issued for, the public key, and the certificate's expiration date. A CA uses one or more signing certificates to create SSL certificates. Each signing certificate has a matching private key that is used to create the CA signature. The CA makes signed certificates (with the public key embedded) readily available, enabling anyone to use them to verify that an SSL certificate was actually signed by a specific CA.

In general, setting up certificates in both High Availability (HA) and non-HA environments involves the following steps:

1. Generating an identity certificate for a server.
2. Installing the identity certificate on the server.
3. Installing the corresponding root certificate on your client or browser.

The specific tasks you need to complete will vary depending on your environment.

Note the following:

- The start-stop sequencing of servers needs to be done carefully in HA environments.
- Non-HA environments, where a virtual IP address is configured, require the completion of a more complicated certificate request process.

1-Way SSL Authentication

This authentication method is used when a client needs assurance that it is connecting to the right server (and not an intermediary server), making it suitable for public resources like online banking websites. Authentication begins when a client requests access to a resource on a server. The server on which the resource resides then sends its server certificate (also known as an SSL certificate) to the client in order to verify its identity. The client then verifies the server certificate against another trusted object: a server root certificate, which must be installed on the client or browser. After the server has been verified, an encrypted (and therefore secure) communication channel is established. At this point, the server prompts for the entry of a valid username and password in an HTML form. Entering user credentials after an SSL connection is established protects them from being intercepted by an unauthorized party. Finally, after the username and password have been accepted, access is granted to the resource residing on the server.

Note A client might need to store multiple server certificates to enable interaction with multiple servers.

To determine whether you need to install a root certificate on your client, look for a lock icon in your browser's URL field. If you see this icon, this generally indicates that the necessary root certificate has already been installed. This is usually the case for server certificates signed by one of the bigger Certifying Authorities (CAs), because root certificates from these CAs are included with popular browsers.

If your client does not recognize the CA that signed a server certificate, it will indicate that the connection is not secure. This is not necessarily a bad thing. It just indicates that the identity of the server you want to connect has not been verified. At this point, you can do one of two things: First, you can install the necessary

root certificate on your client or browser. A lock icon in your browser's URL field will indicate the certificate was installed successfully. And second, you can install a self-signed certificate on your client. Unlike a root certificate, which is signed by a trusted CA, a self-signed certificate is signed by the person or entity that created it. While you can use a self-signed certificate to create an encrypted channel, understand that it carries an inherent amount of risk because the identity of the server you are connected with has not been verified.

Disable Insecure Ports and Services

As a general policy, any ports that are not needed should be disabled. You need to first know which ports are enabled, and then decide which of these ports can be safely disabled without disrupting the normal functioning of . You can do this by listing the ports that are open and comparing it with a list of ports needed for .

To view a list of all open listening ports:

Step 1 Log in as a Linux CLI admin user and enter the **netstat -aln** command.
The **netstat -aln** command displays the server's currently open (enabled) TCP/UDP ports, the status of other services the system is using, and other security-related configuration information. The command returns output similar to the following:

```
[root@vm ~]# netstat -aln
Active Internet connections (servers and established)
Proto  Recv-Q   Send-Q   Local Address           Foreign Address         State
tcp    0        0        0.0.0.0:111             0.0.0.0:*               LISTEN
tcp    0        0        127.0.0.1:8080          0.0.0.0:*               LISTEN
tcp    0        0        0.0.0.0:22              0.0.0.0:*               LISTEN
tcp    0        0        127.0.0.1:25            0.0.0.0:*               LISTEN
tcp    0        0        127.0.0.1:10248         0.0.0.0:*               LISTEN
tcp    0        0        127.0.0.1:10249         0.0.0.0:*               LISTEN
tcp    0        0        192.168.125.114:40764   192.168.125.114:2379    ESTABLISHED
tcp    0        0        192.168.125.114:48714   192.168.125.114:10250   CLOSE_WAIT
tcp    0        0        192.168.125.114:40798   192.168.125.114:2379    ESTABLISHED
tcp    0        0        127.0.0.1:33392         127.0.0.1:8080          TIME_WAIT
tcp    0        0        192.168.125.114:40814   192.168.125.114:2379    ESTABLISHED
tcp    0        0        192.168.125.114:40780   192.168.125.114:2379    ESTABLISHED
tcp    0        0        127.0.0.1:8080          127.0.0.1:44276         ESTABLISHED
tcp    0        0        192.168.125.114:40836   192.168.125.114:2379    ESTABLISHED
tcp    0        0        192.168.125.114:40768   192.168.125.114:2379    ESTABLISHED
tcp    0        0        127.0.0.1:59434         127.0.0.1:8080          ESTABLISHED
tcp    0        0        192.168.125.114:40818   192.168.125.114:2379    ESTABLISHED
tcp    0        0        192.168.125.114:22      192.168.125.1:45837     ESTABLISHED
tcp    0        0        127.0.0.1:8080          127.0.0.1:48174         ESTABLISHED
tcp    0        0        127.0.0.1:49150         127.0.0.1:8080          ESTABLISHED
tcp    0        0        192.168.125.114:40816   192.168.125.114:2379    ESTABLISHED
tcp    0        0        192.168.125.114:55444   192.168.125.114:2379    ESTABLISHED
```

Step 2 Check the *Cisco Crosswork Optimization Engine Installation Guide* for the table of ports used by Cisco Crosswork Optimization Engine, and see if your ports are listed in that table. That table will help you understand which services are using the ports, and which services you do not need—and thus can be safely disabled. In this case, *safe* means you can *safely disable the port without any adverse effects to the product*.

Note If you are not sure whether you should disable a port or service, contact your Cisco representative.

Step 3 If you have firewalls in your network, configure the firewalls to only allow traffic that is needed for Cisco Crosswork Optimization Engine to operate.

Harden Your Storage

We recommend that you secure all storage elements that will participate in your installation, such as the database, backup servers, and so on.

- If you are using external storage, contact your storage vendor and your Cisco representative.
- If you are using internal storage, contact your Cisco representative.
- If you ever uninstall or remove , make sure that all VM-related files that might contain sensitive data are digitally shredded (as opposed to simply deleted). Contact your Cisco representative for more information.

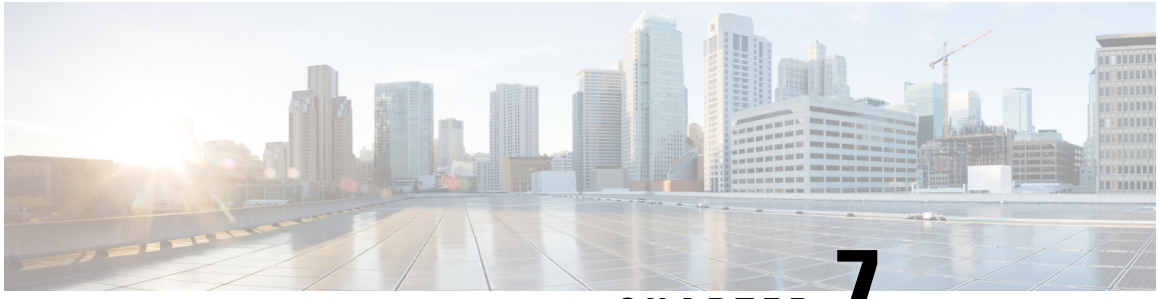

CHAPTER 7

Manage Cisco Crosswork Data Gateway

Networks maintain a large amount of data that spans thousands of devices. Cisco Crosswork Optimization Engine Collection Service collects and manages this data through its integral component - Cisco Crosswork Data Gateway.

This section contains the following topics:

- Overview of Cisco Crosswork Data Gateway, on page 159
- Manage Cisco Crosswork Data Gateway Instances, on page 159
- Configure Cisco Crosswork Data Gateway Settings, on page 175

Overview of Cisco Crosswork Data Gateway

When Cisco Crosswork Optimization Engine and Cisco Crosswork Data Gateway are deployed together, Cisco Crosswork Optimization Engine acts as the **controller application** for the Cisco Crosswork Data Gateway instance. You can use the UI to add and manage additional instances of Cisco Crosswork Data Gateway no matter if they are forwarding data to Cisco Crosswork Optimization Engine or other compatible data consumers. The number of Cisco Crosswork Data Gateway you need depends on the number of devices being supported, the amount of data being processed and your network architecture.

Cisco Crosswork Data Gateway can also be deployed with other Crosswork products and in that case, will have a different controller application.

Note This chapter explains only the Cisco Crosswork Data Gateway features that can be accessed via Cisco Crosswork Optimization Engine UI.

For more information about Cisco Crosswork Data Gateway VM and how to manage it, see **Appendix B:** Configure Cisco Crosswork Data Gateway Base VM, on page 207.

We also recommended that you read about components of Cisco Crosswork Data Gateway at Cisco Crosswork Data Gateway Components, on page 210 before moving further.

Manage Cisco Crosswork Data Gateway Instances

Cisco Crosswork Data Gateway is initially deployed with just a basic VM called the Base VM (containing only enough software to register itself with its controller).

It follows the instructions from Crosswork - collects data as requested and sends it to the defined output destination.

Depending on your private network's size and configuration, you may require one or more Cisco Crosswork Data Gateway instances for collection. It may be necessary to deploy multiple Cisco Crosswork Data Gateway instances to address the requirements for:

1. Geo-separated regions
2. Massive scale

Cisco recommends the simplest approach of a fixed configuration of devices to a particular instance (such as x to y for CDG1 and (y+1) to z for CDG2).

Note More complicated approaches for resource optimization and dynamic assignment of tasks are possible and if desired, we recommend working with Cisco Customer Experience team to design the behavior.

To open Cisco Crosswork Data Gateway management view, choose **Admin** > **Data Gateway Management**.

Figure 57: Data Gateway Management View

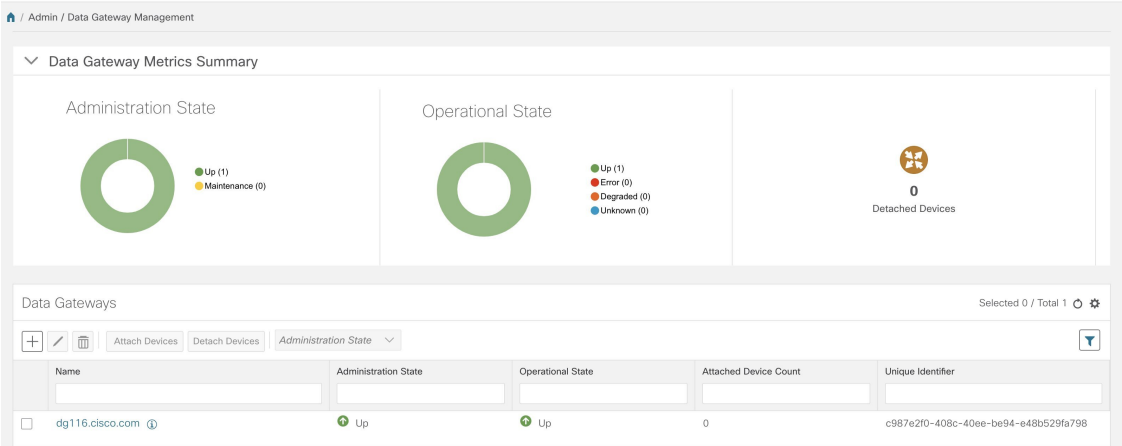

Item	Description
Data Gateway Metrics Summary Pane	Summarizes the overall metrics of all Cisco Crosswork Data Gateway instances currently enrolled with Crosswork: • **Administration State Tile**: shows the number of Cisco Crosswork Data Gateway instances in each adminstration state i.e., Up and Maintenance. • **Operational State Tile**: shows the number of Cisco Crosswork Data Gateway instances in each operational state i.e., Up, Error, Degraded, and Unknown. • **Detached Devices Tile**: Shows the number of devices that are currently not attached to any Cisco Crosswork Data Gateway instance.

Item	Description

Data Gateways Pane	Provides options to add, edit, and delete Cisco Crosswork Data Gateway VMs, attach/detach devices, change administration state, and filter options.
	
	It also displays the following details for the individual Cisco Crosswork Data Gateway instances:
	
	• **Name**: Name of the Cisco Crosswork Data Gateway VM.
	
	• **Administration State**: Adminstration state of the Cisco Crosswork Data Gateway VM. A Cisco Crosswork Data Gateway VM has either of the two states at a time:
	
	  • **Up**: The VM is currently active.
	
	  • **Maintenance**: The VM is not operational ("down") and has been set to "Maintenance" mode by the user. No new jobs are submitted to Cisco Crosswork Data Gateway while it is in this mode. However, the currently running collection jobs do not stop.
	
	• **Operational State**: Operational state of the Cisco Crosswork Data Gateway VM. A Crosswork Data Gateway VM has either of the four states at a time:
	
	  • **Up**: The VM is operational and all individual components are "OK".
	
	  • **Error**:
	
	  The VM's operational state is in an error condition. It is either not reachable or all the critical components on the VM are "not OK".
	
	  • **Degraded**:
	
	  The VM's operational state is degraded as one or more critical components on the VM are "not OK".
	
	  • **Unknown**:
	
	  The VM's operational state is unknown as it has enrolled itself with Crosswork, but hasn't established a session yet.

From the **Data Gateways** pane, you can add a new Cisco Crosswork Data Gateway instance, update the settings configured for an existing instance, de-enroll an instance, attach devices to an instance, detach devices from a instance, or change adminstration state of an instance.

Figure 58: Data Gateways Pane

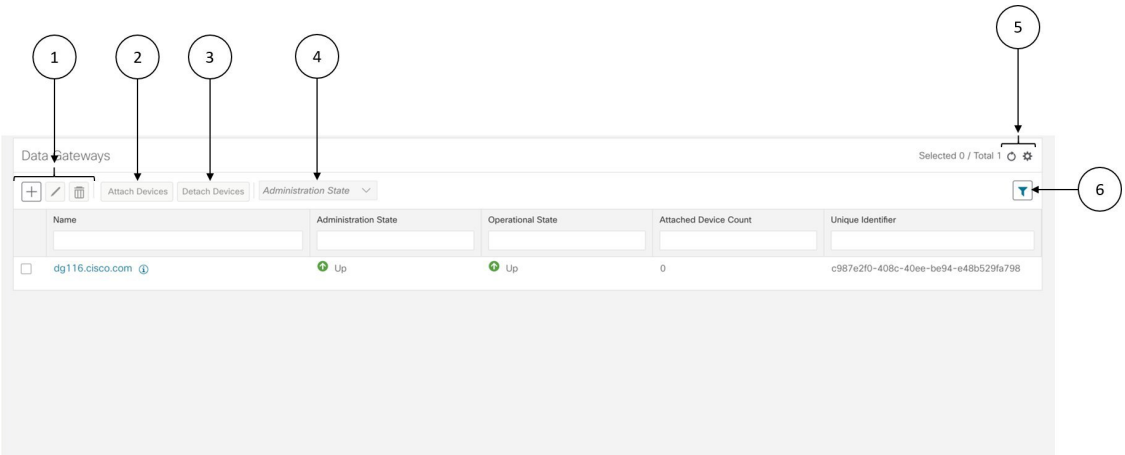

Item	Description
1	Click ⊞ to add a Cisco Crosswork Data Gateway VM. See Add a Cisco Crosswork Data Gateway Instance, on page 164.
	Click ✎ to edit the settings for the selected Cisco Crosswork Data Gateway VM. See Update Cisco Crosswork Data Gateway Instance Enrollment Settings, on page 164.
	Click 🗑 to de-enroll the selected Cisco Crosswork Data Gateway VM. See De-enroll a Cisco Crosswork Data Gateway Instance, on page 168.
2	Click **Attach Devices** to attach devices to the selected Cisco Crosswork Data Gateway VM. See Attach a Device to a Cisco Crosswork Data Gateway Instance, on page 169.
3	Click **Detach Devices** to detach devices from the selected Cisco Crosswork Data Gateway VM. See Detach a Device From a Cisco Crosswork Data Gateway Instance, on page 171.
4	Click **Adminstration State** to switch administration state of the selected Data Gateway VM. See Change the Administration State of a Cisco Crosswork Data Gateway Instance, on page 167.
5	Click ↻ to refresh the **Data Gateways** window.
	Click ⚙ to choose the columns to make visible in the **Data Gateways** window (see Set, Sort and Filter Table Data, on page 7).
6	Click ▼ to set filter criteria on one or more columns in the **Data Gateways** window.
	Click the **Clear All Filters** link to clear any filter criteria you may have set.

The **Data Gateways** pane displays the following details of the enrolled Cisco Crosswork Data Gateway instances:

Field	Description
Name	Name of the Cisco Crosswork Data Gateway.
Administration State	Adminstration state of the Cisco Crosswork Data Gateway instance.
Operational State	Operational state of the Cisco Crosswork Data Gateway instance.
Attached Device Count	Number of devices attached to the Cisco Crosswork Data Gateway instance.
Unique Identifier	Unique identifier of the Cisco Crosswork Data Gateway instance.

Add a Cisco Crosswork Data Gateway Instance

After installing Cisco Crosswork Data Gateway, you must enroll it with Cisco Crosswork Optimization Engine.

Steps to enroll a Cisco Crosswork Data Gateway instance is described in *Cisco Crosswork Optimization Engine Installation Guide* in Section: **Enroll Cisco Crosswork Data Gateway With Cisco Crosswork Optimization Engine**

After enrolling, you must verify that the operational state of the Cisco Crosswork Data Gateway instance is **Up** before beginning to use it.

Note Watch out for "alerts" at the top of the **Data Gateway** page while the Cisco Crosswork Data Gateway is not operationally up.

Update Cisco Crosswork Data Gateway Instance Enrollment Settings

If there's an update for the Cisco Crosswork Data Gateway VM, you can regenerate a new enrollment package and upload it to Crosswork by following these steps:

Before you begin

Ensure that you have manually copied the new enrollment package to your local PC as per the procedure described in the *Cisco Crosswork Optimization Engine Installation Guide* in Section: *Export Enrollment Package*.

Step 1 From the main menu, choose **Admin** > **Data Gateway Management**. The **Data Gateway Management** view opens.

Step 2 From the **Data Gateways** window, select the Cisco Crosswork Data Gateway instance you want to update.

Step 3 Click to edit the settings for the selected Cisco Crosswork Data Gateway instance.

Step 4 In the **Update Data Gateway** pop up, click **Browse** to select the new enrollment package.

Select the admin state in which you want to bring up the Cisco Crosswork Data Gateway instance.

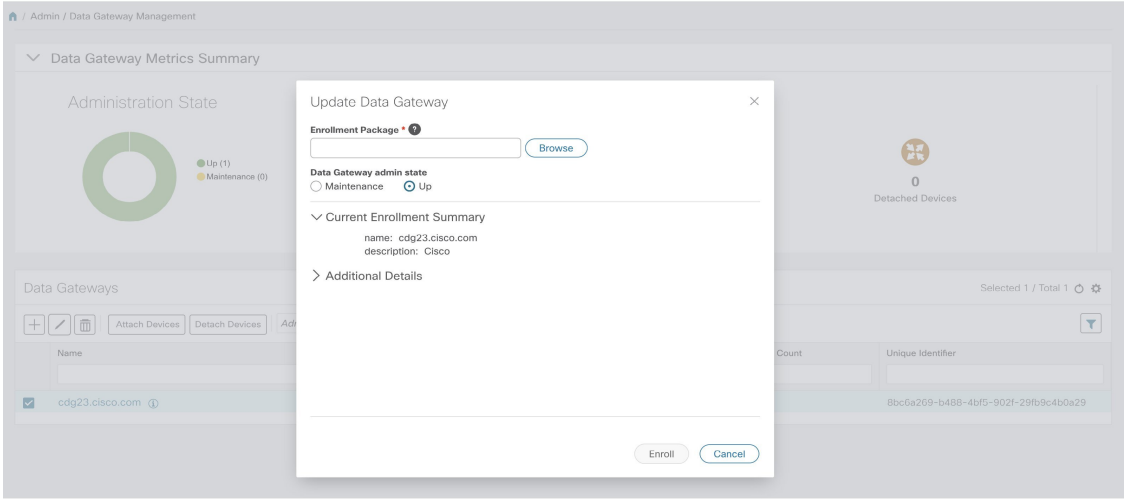

Step 5 Click **Enroll**.

View Enrollment Details

To view enrollment details of a Cisco Crosswork Data Gateway instance, in the **Data Gateways** pane, click icon next to the Cisco Crosswork Data Gateway name as shown in the following figure.

Note Some of these details are the OVF parameters that were configured in the OVA Template while installing Cisco Crosswork Data Gateway. For description of these parameters, see Section: **Install Crosswork Data Gateway** in *Cisco Crosswork Optimization Engine Installation Guide*.

Figure 59: Crosswork Data Gateway Enrollment Details

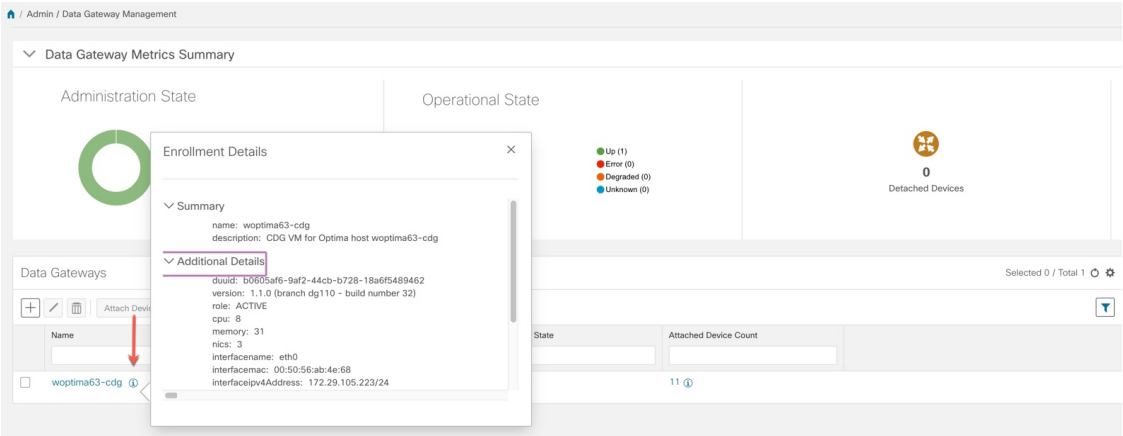

Following enrollment details are displayed:

Field	Description
Summary	
name	Name of the Cisco Crosswork Data Gateway instance.
description	User-friendly description to be displayed in the controller i.e., Crosswork.
Additional Details	
duuid	Unique identifier for the Cisco Crosswork Data Gateway instance.
version	Currently installed version of Cisco Crosswork Data Gateway.
role	Is the Cisco Crosswork Data Gateway instance active or in maintenance mode.
cpu	Number of vCPUs.
memory	Amount of total memory.
nics	Number of NICs being used by Cisco Crosswork Data Gateway. This is 3 in case of on-premise installation i.e., for Cisco Crosswork Optimization Engine.
interfacename	Name of the interface.
interfacemac	MAC address of the interface
interfaceIPv4address/interfaceIPv6address	IPv4/IPv6 address of the interface.

Field	Description
cert_chain	Certificate used for handshake between Cisco Crosswork Data Gateway instance and Cisco Crosswork Optimization Engine.

Change the Administration State of a Cisco Crosswork Data Gateway Instance

You can change the adminstration state of a Cisco Crosswork Data Gateway instance via Crosswork UI.

Note If the maintenance activities are affecting the communication between Crosswork and Cisco Crosswork Data Gateway, the collection is interrupted and resumes when the communication is restored.

While an instance is in in **Maintenance** mode, no new jobs are submitted to it. During downtime, admin can do modifications to Cisco Crosswork Data Gateway, such as updating the certificates, changing management address, etc.

Once changes are done, Admin can change the administration state to **Up**. Once the Cisco Crosswork Data Gateway is up, Crosswork resumes sending jobs to it.

Follow the steps below to change the administration state of a Cisco Crosswork Data Gateway instance.

Step 1 From the main menu, choose **Admin** > **Data Gateway Management**. The **Data Gateway Management** view opens.

Step 2 From the **Data Gateways** window, select the Cisco Crosswork Data Gateway instance whose administration state you want to change.

Step 3 From the **Administration State** dropdown, select the state to which you want to switch to.

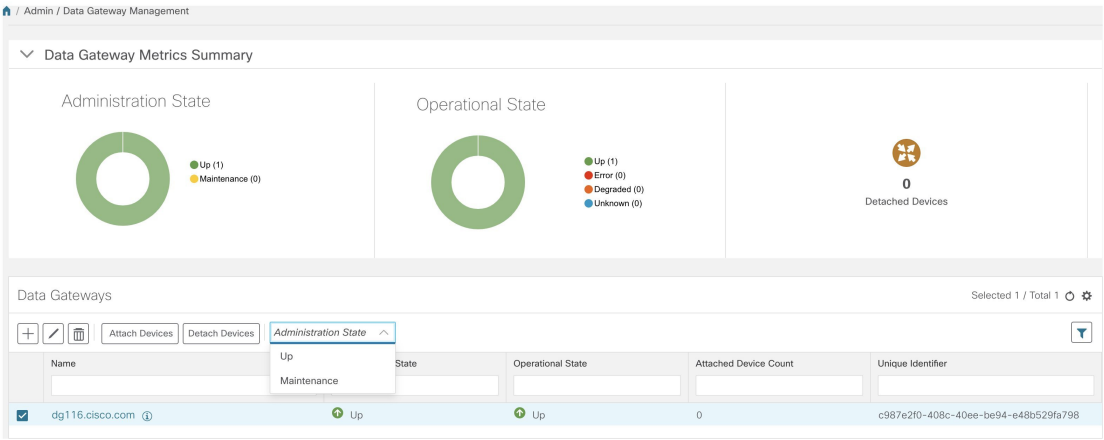

De-enroll a Cisco Crosswork Data Gateway Instance

Follow the steps below to de-enroll a Cisco Crosswork Data Gateway instance.

Step 1 From the main menu, choose **Admin** > **Data Gateway Management**. The **Data Gateway Management** view opens.

Step 2 From the **Data Gateways** window, select the Cisco Crosswork Data Gateway instance you want to delete.

Step 3 Click 🗑.

Step 4 A Cisco Crosswork Data Gateway instance must be in maintenance mode to be deleted. Click **Switch & Continue** when prompted to switch to **Maintenance** mode.

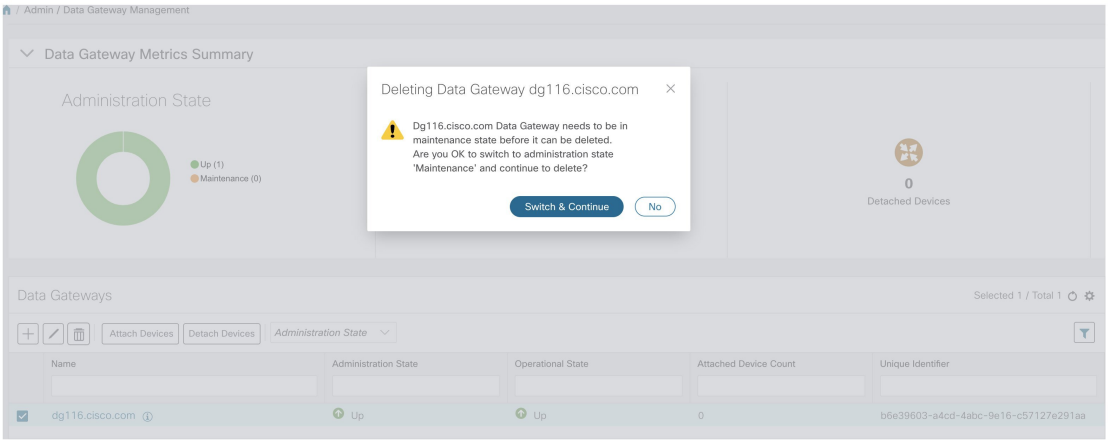

The selected Cisco Crosswork Data Gateway instance is deleted.

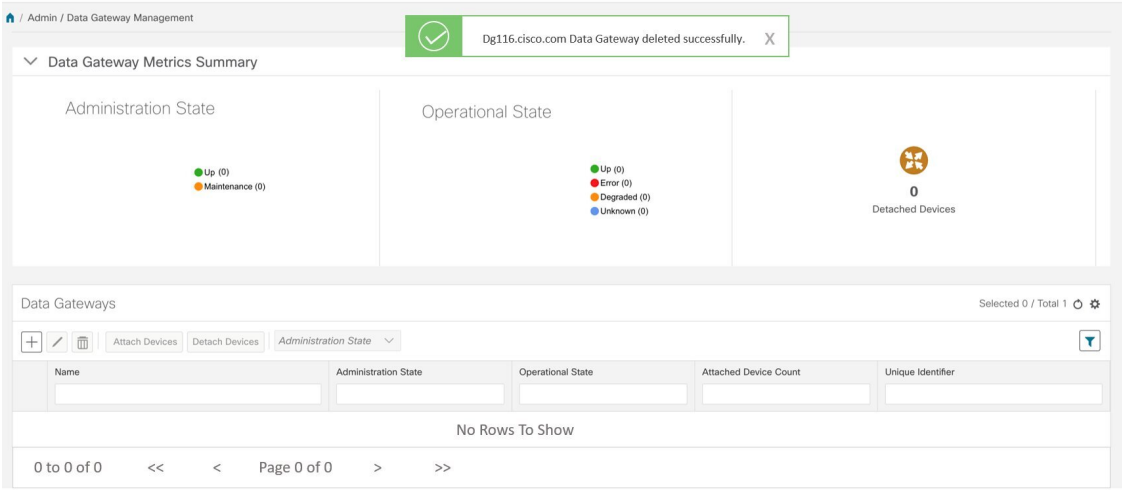

Attach a Device to a Cisco Crosswork Data Gateway Instance

Note A device can only be attached to one Cisco Crosswork Data Gateway instance.

Follow the steps below to attach a device to a Cisco Crosswork Data Gateway instance.

Before you begin

1. For optimal performance, it is recommended that device attaching to Cisco Crosswork Data Gateway instance should be done in batches of no more than 300 devices.

   You can add more than 300 devices. However, doing so may cause a performance impact.

2. Ensure that both the adminstration state and operational state of the Cisco Crosswork Data Gateway instance to which you want to attach devices is "Up". Only then proceed with attaching devices.

Step 1 From the main menu, choose **Admin** > **Data Gateway Management**. The **Data Gateway Management** view opens.

Step 2 From the **Data Gateways** window, select the Cisco Crosswork Data Gateway instance to which you want to attach devices.

Step 3 Click **Attach Devices**. The **Attach Devices** window opens. It lists all the devices available for attaching.

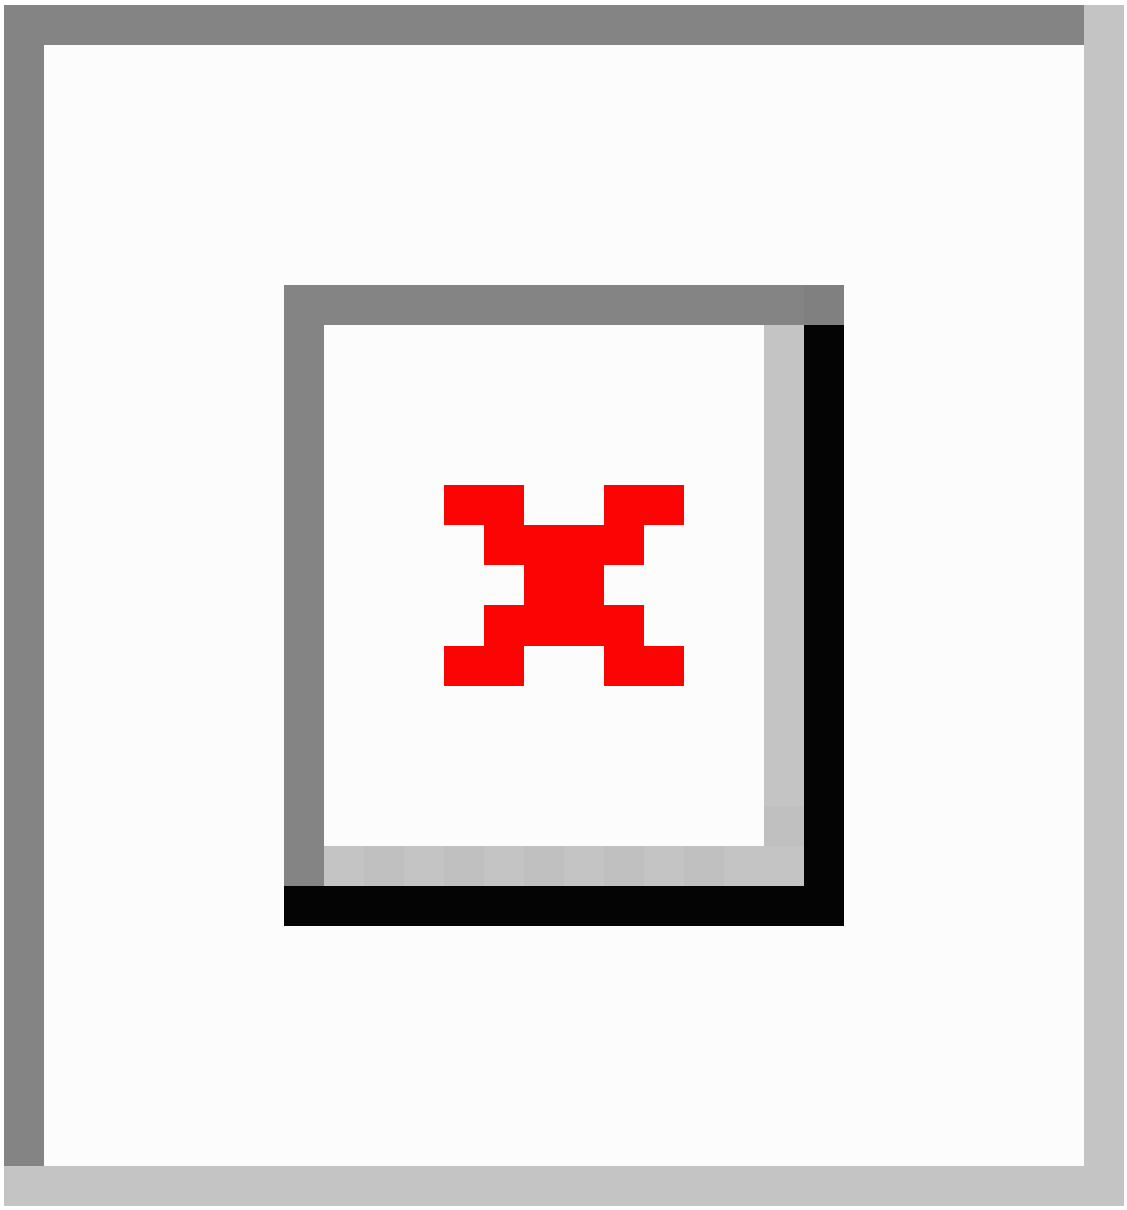

Step 4 To attach all the devices, click **Attach All Devices**. Otherwise, select the devices you want to attach and click **Attach Selected Devices**.

What to do next

To verify if the devices were attached to the VM, check the **Attached Device Count** under the **Data Gateways** pane. The count would have increased.

Click on the icon next to the attached device count to see the list of all devices attached to the selected instance, as shown in the following figure.

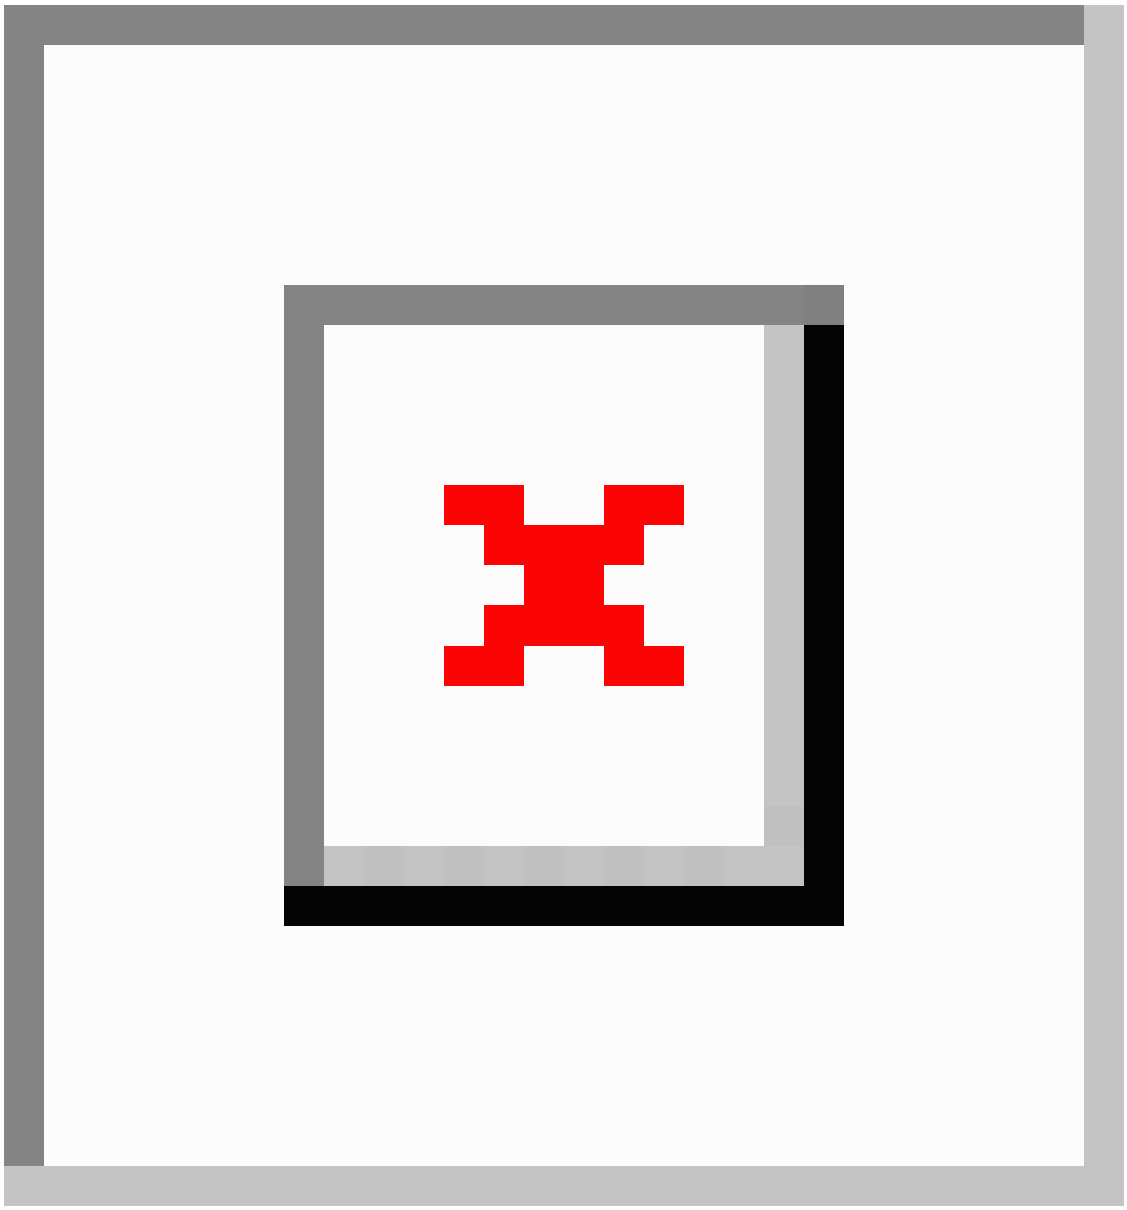

Detach a Device From a Cisco Crosswork Data Gateway Instance

Follow the steps below to detach a device from a Cisco Crosswork Data Gateway instance.

Step 1 From the main menu, choose **Admin** > **Data Gateway Management**. The **Data Gateway Management** view opens.

Step 2 From the **Data Gateways** window, select the Cisco Crosswork Data Gateway instance from which you want to detach devices.

Step 3 Click **Detach Devices**. The **Detach Devices** window opens. It lists all the devices attached to the selected Cisco Crosswork Data Gateway instance.

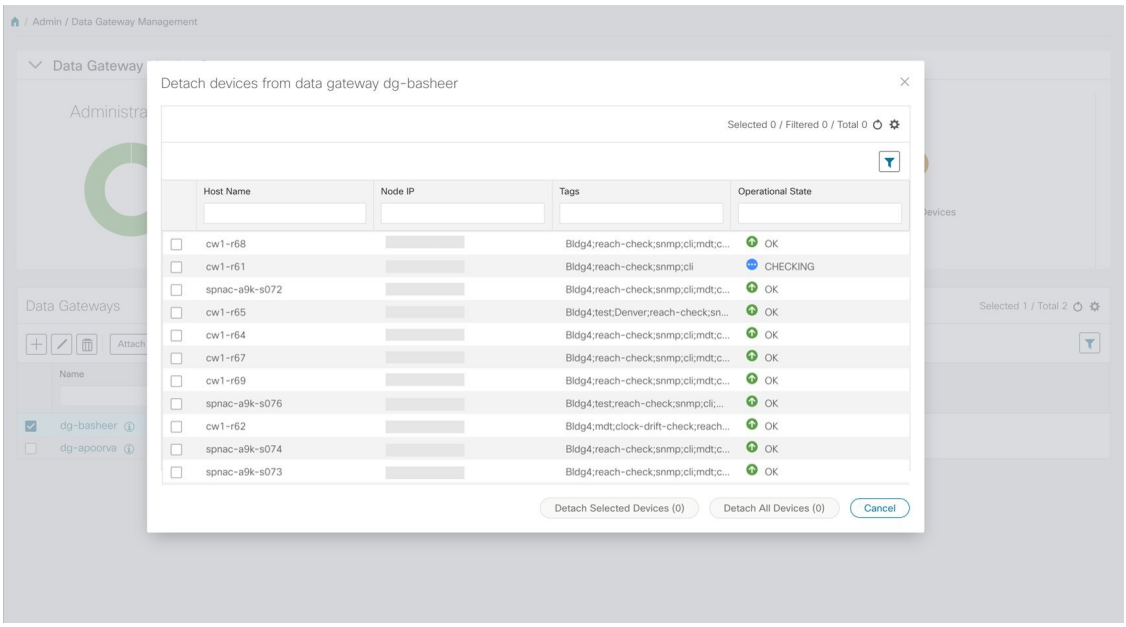

Step 4 To detach all the devices click **Detach All Devices**. Otherwise, select the devices you want to detach and click **Detach Selected Devices**.

What to do next

To verify if the devices were detached from the VM, check the **Attached Device Count** under **Data Gateways** window. The count would have decreased.

View Cisco Crosswork Data Gateway Instance Health

Cisco Crosswork Data Gateway comprises of various containerized services running on an Ubuntu VM. Its overall health depends on health of each containerized service.

Cisco Crosswork Data Gateway collects host and container metrics and writes them to a container mounted path in vitals.json file and sends it to the Crosswork.

Vitals also contains the health information of individual container services running on the Cisco Crosswork Data Gateway instance and their resource consumption.

To view health of a Cisco Crosswork Data Gateway instance, in the **Data Gateways** window, click the name of the Cisco Crosswork Data Gateway instance whose health you want to view as shown in the following figure.

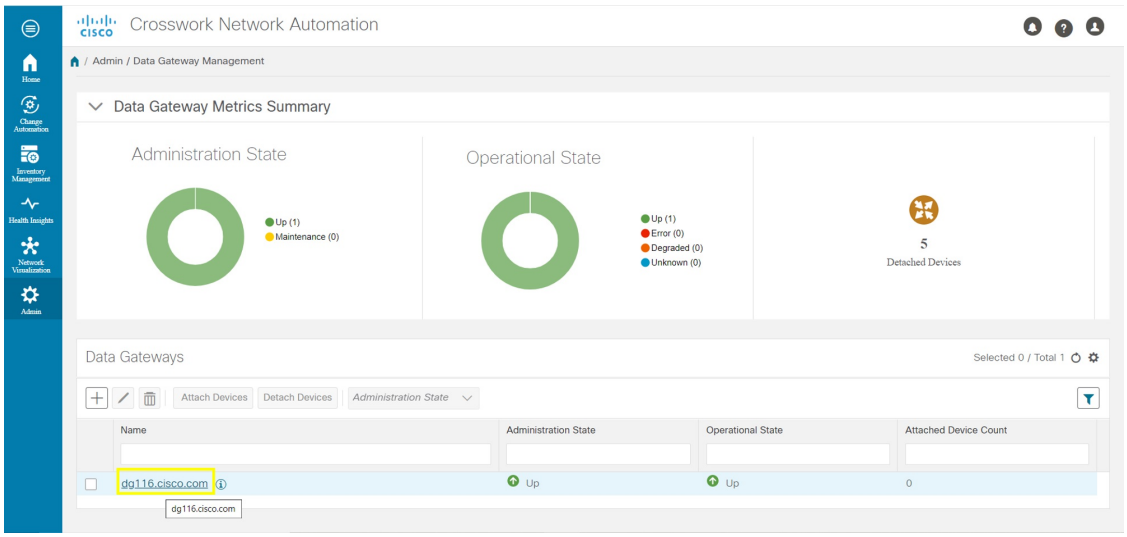

The **Health** pop up displays the following details:

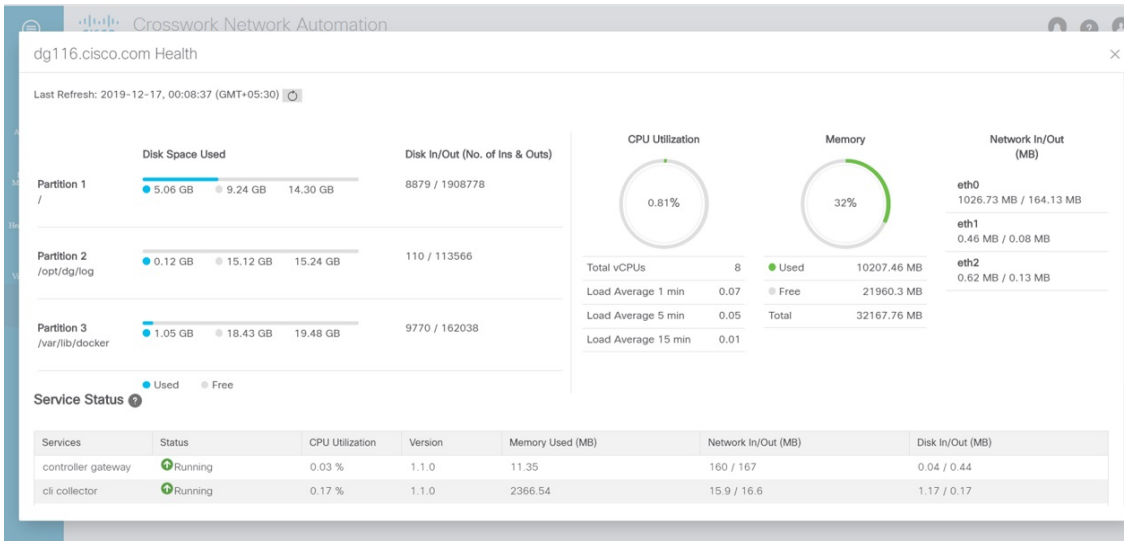

Field	Description
Host VM	
Last Refresh	Date and time of the last refresh. Click ⟳ to refresh the **Data Gateway Health** pop up.
Disk Space Used	Percentage of the disk space used for partitions: / /opt/dg/log /var/lib/docker

Field	Description
Disk In/Out	Number of read/write or input/output operations involving a disk for the partitions: / /opt/dg/log /var/lib/docker **Note** This is a cumulative counter, not a delta time series.
CPU Utilization	Amount of actively used CPU and total number of vCPUs.
Load	Load average – is the average system load over a given period of time of 1, 5, and 15 minutes.
Memory	Amount of memory used and available memory.
Network In/Out	The amount of data sent/received in MB for NIC interfaces: eth0 eth1 eth2 **Note** This is a cumulative counter, not a delta time series.
Service Status	
Service	Name of the Cisco Crosswork Data Gateway service.
Status	Status of the service: • Running • Degraded • Error
CPU Utilization	Percentage of actively utilized CPU by the service.
Version	Version of the service deployed.
Memory Used (MB)	Amount of memory being used by the service.
Network In/Out	The amount of data sent/received in MB by the service over its interface. **Note** This is a cumulative counter, not a delta time series.

Field	Description
Disk In/Out	Number of read/write or input/output operations that the service has done involving a disk. **Note** This is a cumulative counter, not a delta time series.

Configure Cisco Crosswork Data Gateway Settings

This section describes how to configure global settings for Cisco Crosswork Data Gateway i.e., managing data destinations and custom software packages.

To open Cisco Crosswork Data Gateway global settings view, choose **Admin** > **Data Gateway Global Settings** from the left navigation bar in the Cisco Crosswork Optimization Engine window.

Figure 60: Data Gateway Global Settings View

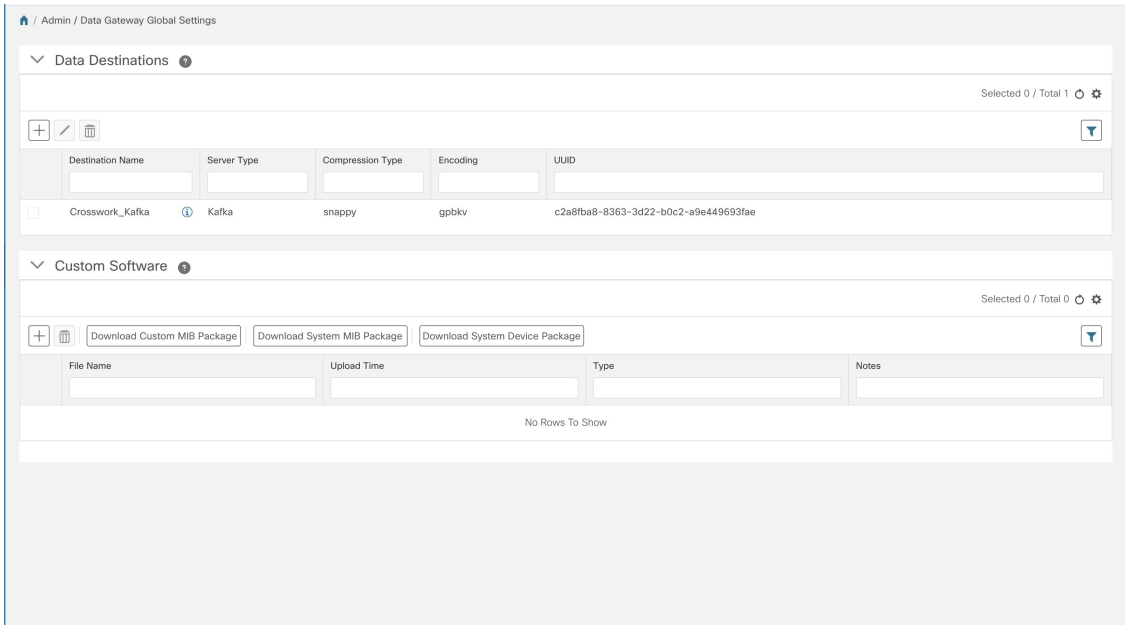

Item	Description
Data Destinations Pane	Shows approved external data destinations that can be used by collection jobs to deposit their data and provides options to add, edit, and delete data destinations.
Custom Software Pane	Provides options to: • add and delete custom MIBs and device packages • download custom MIBs, system MIBs, and device packages

Manage Data Destinations

From the **Data Destinations** pane, you can add a new data destination, update the settings configured for an existing data destination, and delete a data destination.

Note The **Crosswork_Kafka** data destination in the below figure is Cisco Crosswork Optimization Engine's internal data destination and hence, it cannot be updated or deleted.

Figure 61: Data Destinations Pane

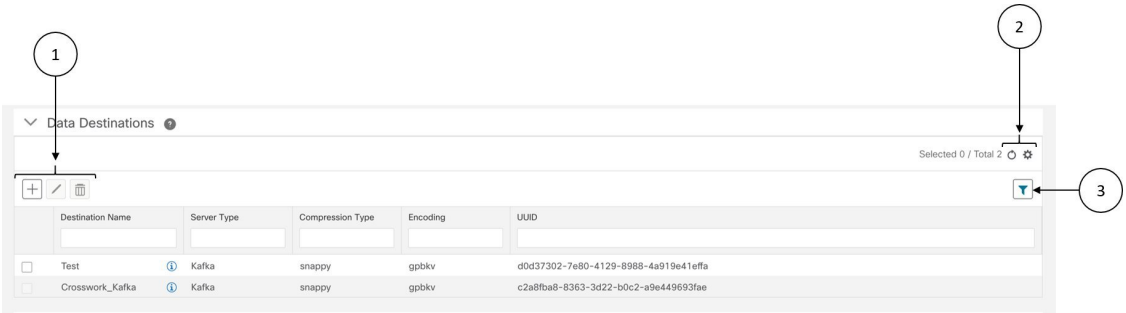

Item	Description
1	Click to add a data destination. See Add a Data Destination, on page 177.
	Click to edit the settings for the selected data destination. See Update a Data Destination, on page 180.
	Click to delete the selected data destination. See Delete a Data Destination, on page 181.
2	Click to refresh the **Data Destinations** window.
	Click to choose the columns to make visible in the **Data Destinations** window (see Set, Sort and Filter Table Data, on page 7).
3	Click to set filter criteria on one or more columns in the **Data Destinations** window.
	Click the **Clear All Filters** link to clear any filter criteria you may have set.

Data Destination pane displays the following details of the data destinations:

Field	Description
Destination Name	Name of the data destination
Server Type	Server type of the data destination i.e., external Kafka or gRPC server.

Field	Description
Compression Type	Compression type being used for the data destination. Crosswork
Encoding	Encoding type being used for the data destination.
UUID	Unique identifier for the data destination. This ID is automatically generated by Crosswork when an external data destination is created and is a required parameter for collection job creation.

Add a Data Destination

Note

- If you reinstall an already existing external Kafka data destination with the same IP address, then the collectors need to be restarted for changes to take place .

Follow the steps below to add a new data destination. You can then use this data destination for data collection. You can also add multiple data destinations.

Before you begin

If you are using an external Kafka server for data collection, ensure the following:

- You have configured the following properties on the external Kafka server:

  Note Refer your Kafka documentation for description and usage of these properties as this explanation is out of scope of this document.

  - `num.io.threads = 8`
  - `num.network.threads = 3`
  - `message.max.bytes= 30000000`

- Create Kafka topics that you want to be used for data collection.

Step 1 From the main menu, choose **Admin**.

Step 2 From **Data Destinations** pane, choose [+].

Step 3 In the **Add Destination** pop-up, enter the values for the following fields as per the table below:

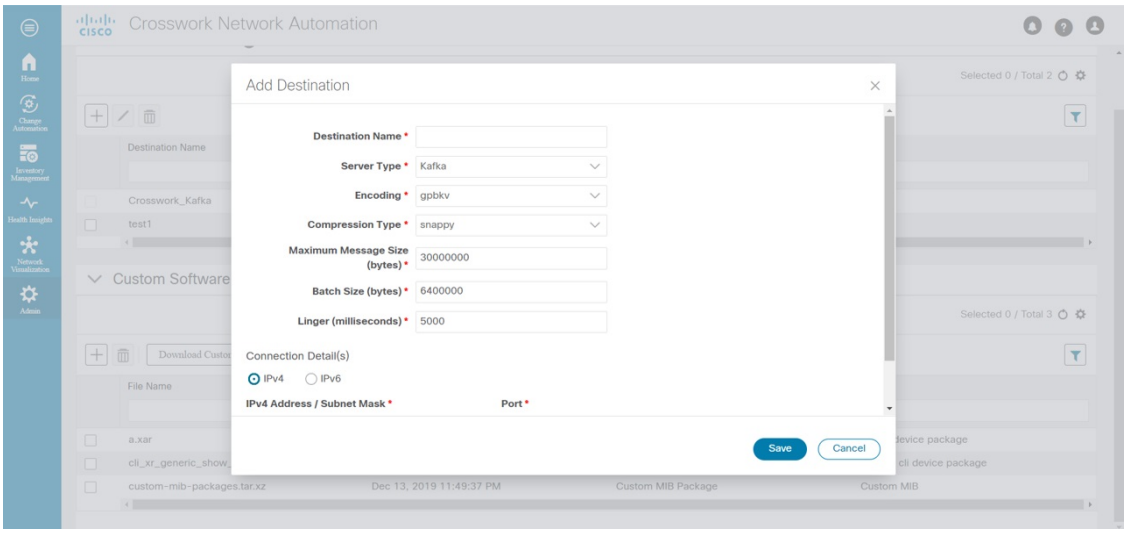

Field	Value
Destination Name	Enter a descriptive data destination name. The name can contain a maximum of 128 alphanumeric characters, plus underscores ("_") or hyphens ("-"). No other special characters are allowed. If you will have many data destinations, make the name as informative as possible to be able to distinguish later.
Server Type	From the drop down, select the server type of your data destination (Kafka/gRPC).
Encoding	From the drop down, select the encoding (json/gpbkv).
Compression Type	From the drop down, select the compression type: Compression types supported for Kafka are snappy, gzip, lz4, zstd, and none) **Note** zstd compression type is supported only for Kafka 2.0 or higher. Compression types supported for gRPC are snappy, gzip, and deflate.
Maximum Message Size (bytes) (Kafka-only)	Enter the maximum message size in bytes. • **Default Value**: 30000000 bytes/ 30 MB • **Min**: 1000000 bytes/1 MB • **Max**: 30000000 bytes/ 30 MB For `Maximum Message Size` property, you can input a value lesser than the default, but not more.

Field	Value
Batch Size (bytes) (Kafka-only)	Enter the required batch size in bytes. • **Default Value**: 6400000 bytes/6.4 MB • **Min**: 16384 bytes/ 16.38 KB • **Max**: 6400000 bytes/6.4 MB **Note** For `Batch Size` property, you can input a value lesser than the default, but not more.
Linger (milliseconds) (Kafka-only)	Enter the required linger time in milliseconds. • **Default Value**: 5000 ms • **Min**: 0 ms • **Max**: 5000 ms

Step 4 Select a protocol from the **Connection Details** options. Cisco Crosswork Data Gateway supports both IPv4 and IPv6.

Step 5 Complete the **Connection Details** fields as described in the following table. The fields displayed will vary with the connectivity type you chose. The values you enter must match the values configured on the device.

Connectivity Type	Fields
IPv4	Enter the required **IPv4 Address/ Subnet Mask**, and **Port**. You can add multiple IPv4 addresses by clicking **+ Add Another**
IPv6	Enter the required **IPv6 Address/ Subnet Mask**, and **Port**. You can add multiple IPv6 addresses by clicking **+ Add Another**.

Step 6 Click **Save**.

What to do next

Create the Kafka topics prior to submitting the job to Crosswork. Depending on external Kafka and how topics are managed in that external Kafka, Cisco Crosswork Data Gateway logs may show the exception listed when and if the topic does not exist at the time of dispatching the collected data to that specific external Kafka / topic. This could be either due to the topic is not yet created or topic got deleted prior to the completion of the requested collection job and dispatching the collected data.

```
destinationContext: topicmdt4
org.apache.kafka.common.errors.UnknownTopicOrPartitionException: This server does not host
 this topic-partition.
```

Update a Data Destination

Note Updating a data destination causes the Cisco Crosswork Data Gateway instance using it to re-establish a session with that data destination. Thus, the data collection is paused and resumes once the session is re-established.

Follow the steps below to update a data destination.

Step 1 From **Data Destinations** window, select the destination you want to update.

Step 2 Click .

Step 3 In the **Edit Destination** pop up, make the required changes.

Note In **Edit** mode, you cannot update **Destination Name** and **Server Type**.

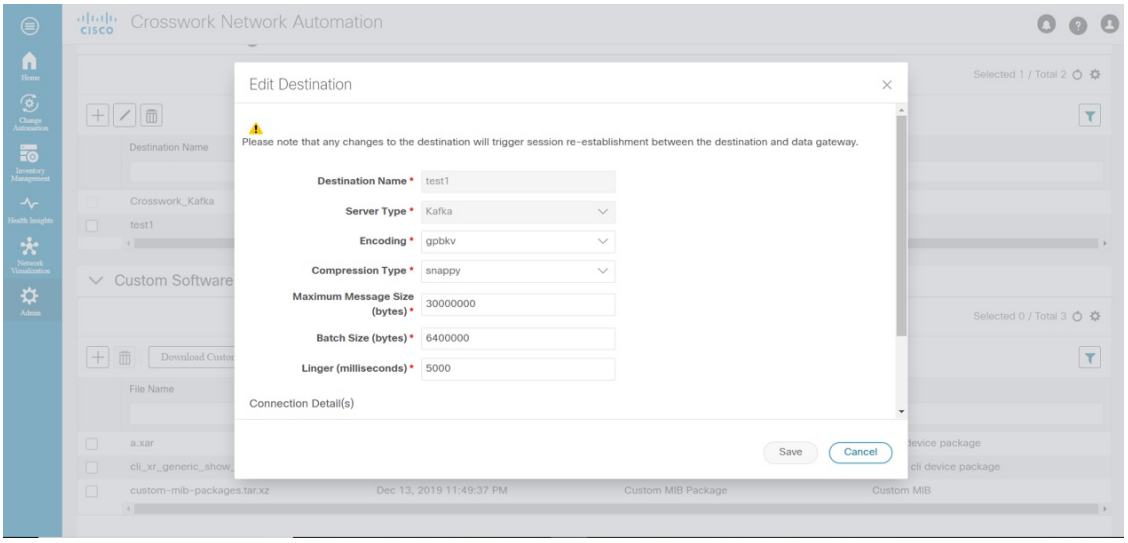

Step 4 Click **Save**.

View Data Destination Details

To view details of a data destination, in the **Data Destinations** pane, click icon next to the data destination name whose details you want to see. Cisco Crosswork Data Gateway displays the details as shown in the following figure.

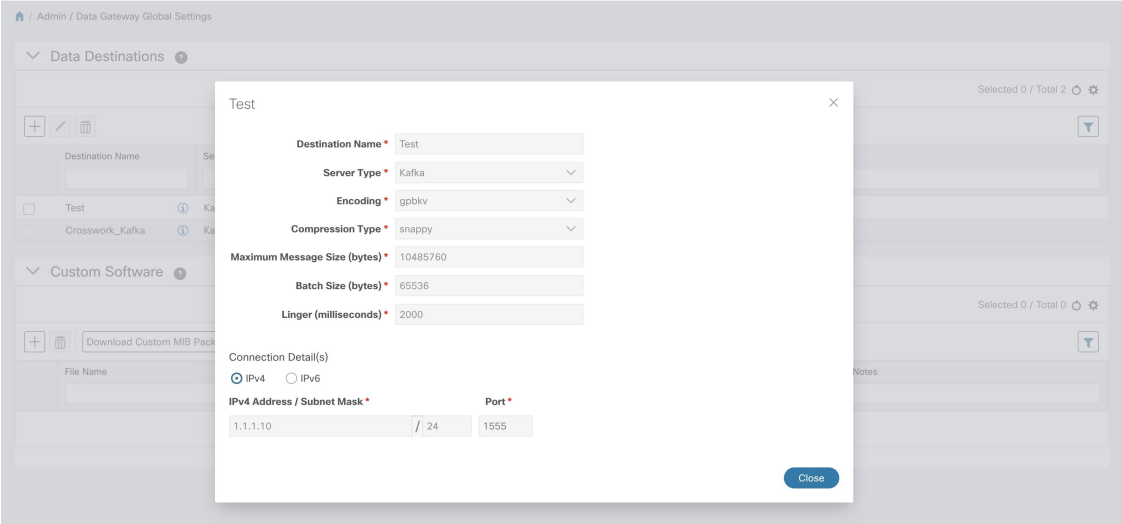

Delete a Data Destination

Follow the steps below to delete a data destination.

Before you begin

A data destination can only be deleted if it is not associated with any collection job. We recommend to check in the **Collection Jobs** view to see if any collection jobs are using the data destination. See Monitoring Collection Jobs, on page 201.

Step 1 From the main menu, choose **Admin**.

Step 2 From the **Data Destinations** pane, select the Data destination(s) you want to delete.

Step 3 Click 🗑.

Step 4 In **Delete Data Destination(s)** pop up, click **Delete** to confirm.

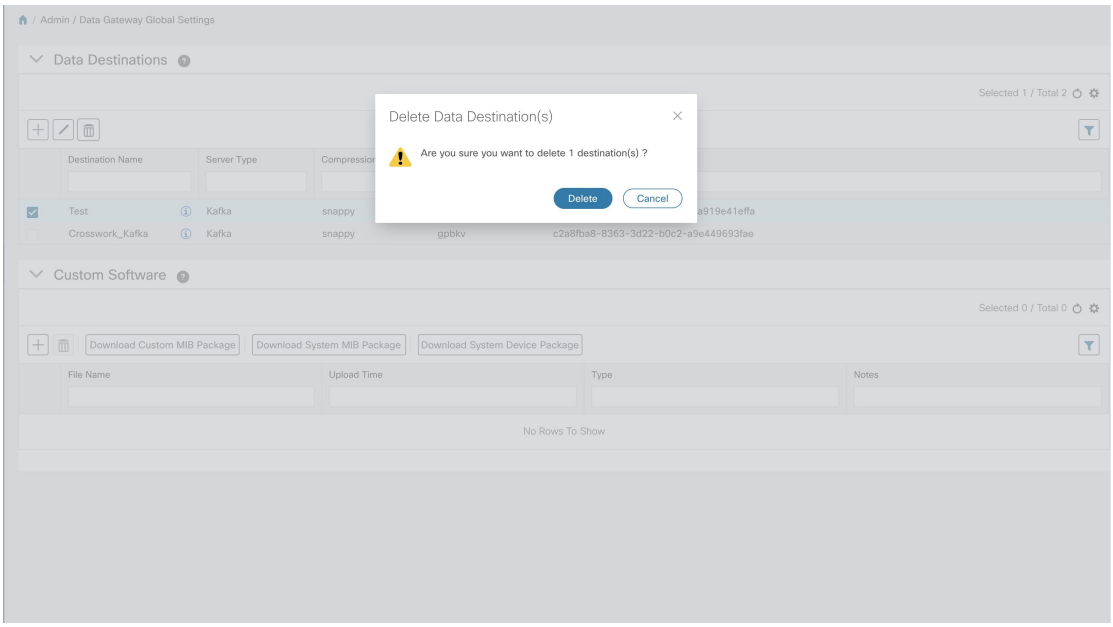

Manage Custom Software Packages

To support third party device CLI and SNMP MIBs, Cisco Crosswork Data Gateway allows you to import the device packages and MIBs to the collectors. Device packages can be imported to allow Cisco Crosswork Data Gateway to retrieve CLI and SNMP data and convert it into xml for third party devices. You can extend the SNMP coverage of Cisco Crosswork Optimization Engine by uploading Custom MIB Packages with any additional MIB and YANG descriptions you require. If you only wish raw SNMP data, no additional files are needed, the system will fold the entire data package into the the Cisco Crosswork Data Gateway data payload.

Note MIBs are required only if the collection request references MIB TABLE names or SCALAR names. However, if the requests are OID-based, then MIBs are not required.

Cisco Crosswork Data Gateway allows you to register and deploy three types of custom software packages:

1. **CLI Device Package**: provides instructions for how to speak to a device using CLI and parse the results into the desired xml.
2. **Custom MIB Packages**: Custom MIBs and device packages can be specific to third party devices or be used to filter the collected data or format it differently for Cisco devices. These are editable by the user.
3. **SNMP Device Package**: provides instructions for how to speak to a device using SNMP and parse the results into the desired xml.

Cisco Crosswork Data Gateway also allows you to download Custom MIB package, System MIB package, and System Device package.

System Device and MIB Packages are bundled in the Crosswork software and are automatically downloaded to the Cisco Crosswork Data Gateway instances. These are NOT modifiable by the user. Custom Device Packages can be downloaded when required for interfacing with third-party devices.

From the **Custom Software** pane, you can add a new custom package, delete a custom package, and download custom packages.

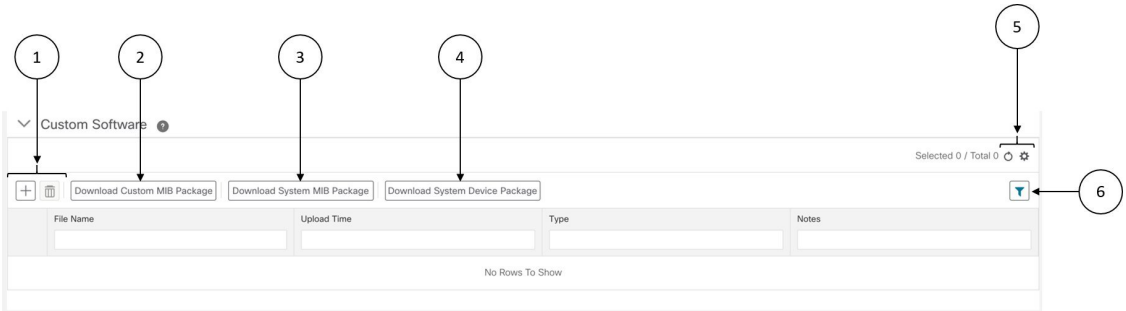

Item	Description
1	Click to add a new custom package. See Add a Custom Software Package, on page 184.
	Click to delete a custom package. See Delete a Custom Software Package, on page 185.
2	Click **Download Custom MIB Package** to download custom MIB packages. See Download Custom or System MIBs and Packages, on page 186.
3	Click **Download System MIB Package** to download system MIB packages. See Download Custom or System MIBs and Packages, on page 186.
4	Click **Download System Device Package** to download system device packages. See Download Custom or System MIBs and Packages, on page 186.
5	Click to refresh the **Custom Software** window.
	Click to choose the columns to make visible in the **Custom Software** window (see Set, Sort and Filter Table Data, on page 7).
6	Click to set filter criteria on one or more columns in the **Custom Software** window.
	Click the **Clear All Filters** link to clear any filter criteria you may have set.

Custom Software pane displays the following details for the available custom software packages:

Field	Description
File Name	Name of the custom software package.
Upload Time	Time of the file upload.
Type	Type of the custom software package.

Field	Description
Notes	Notes related to the custom software package entered by the user while importing the package.

Add a Custom Software Package

Crosswork allows you to upload Custom Device Packages in case you want to filter/format the collected raw data differently.

There are two types of upload:

1. Custom MIB Package upload (a single file custom-mib-packages.tar.xz): which is archive of all custom MIBs/YANGs file
2. Individual Device Package Upload

When uploading new MIBs as a part of Custom MIB Package, it's required that those new MIBs files are loadable within collectors along with existing System MIB files i.e., all dependencies in the files get resolved properly. An offline tool steps are provided for you to ensure that their new MIBs gets parsed and uploaded properly. Accordingly, you can prepare the Custom MIB Package and upload.

For information on how to validate custom MIBs and Yangs i.e., to check if they can be uploaded to Crosswork, see Use Custom MIBs and Yangs on Cisco DevNet.

Note Crosswork doesn't allow Custom MIB package files to overwrite the System MIB Package files. It results in a failed upload attempt.

Using UI, Admin can upload CLI device packages, custom MIB packages, and SNMP device packages. This gets downloaded on the Cisco Crosswork Data Gateway instance to mounted path of respective collectors.

Follow these steps to import a custom software package into Cisco Crosswork Data Gateway:

Note

- Ensure that the custom software package TAR file has just the device package folders and none of the parent folder or hierarchy of folders as part of the TAR file. If not imported properly, Cisco Crosswork Data Gateway throws exceptions when executing the job with custom device package.
- Crosswork does not implement any control on the files being uploaded other than checking the file extension.

Step 1 From the main menu, choose **Admin**.

Step 2 From **Custom Software** window, choose [+].

Step 3 From the **Add Custom Software** pop up, select the type of custom software package you want to import from the **Type** dropdown.

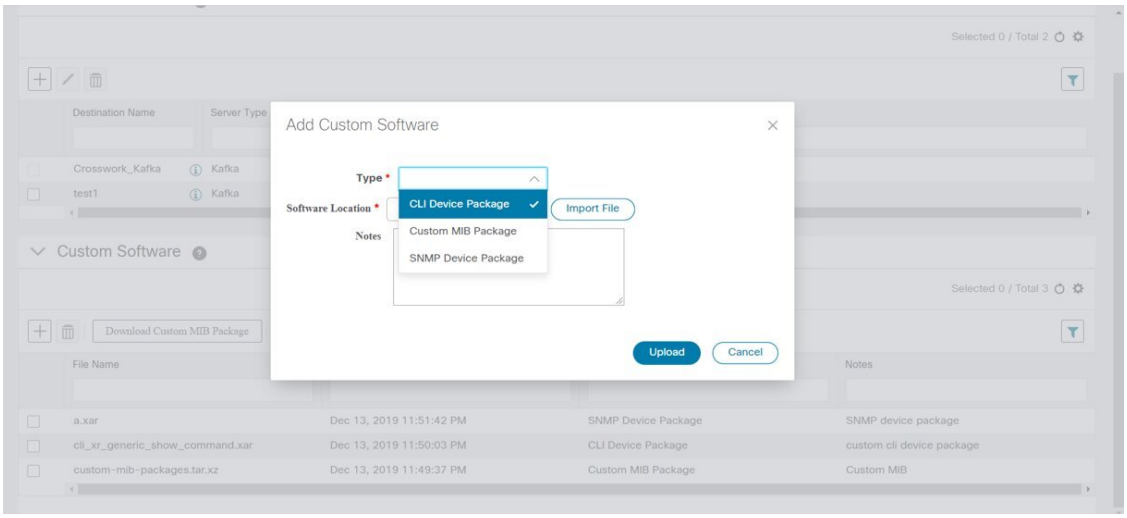

Step 4 Click in the blank field of **Software Location** to open the file browser window and select the custom software package to import and click **Import File**.

Step 5 Add a description of the custom software package in the **Notes** field. This is recommended if you have many packages, to be able to distinguish among them.

Step 6 Click **Upload**.

Delete a Custom Software Package

Follow the steps below to delete a custom software package.

Step 1 From the main menu, choose **Admin**.

Step 2 From the **Custom Software** pane, select the custom package you want to delete.

Step 3 Click 🗑.

Step 4 In the **Delete Custom Software** pop up, click **Delete** to confirm.

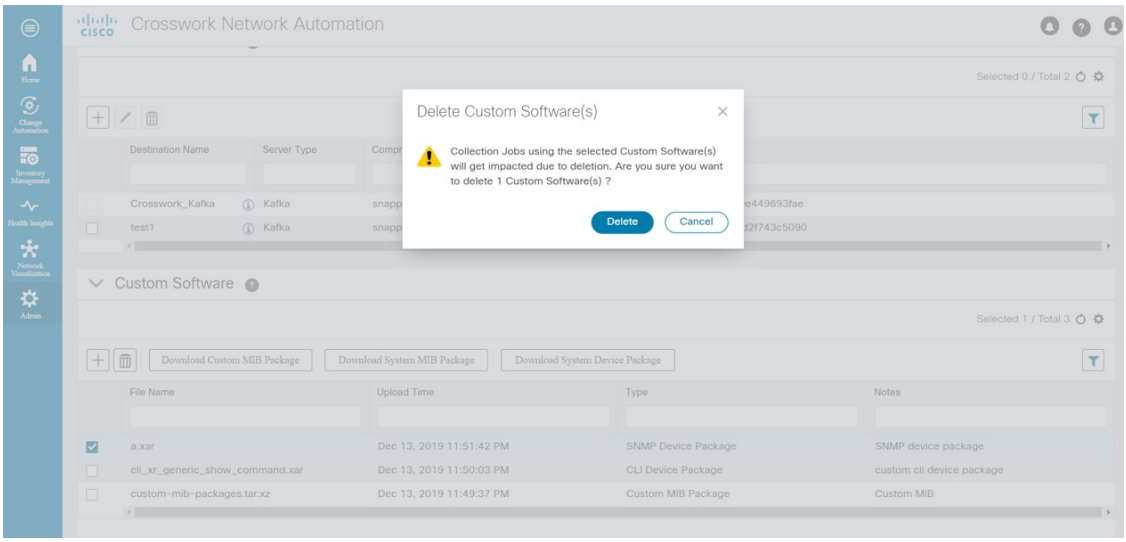

Download Custom or System MIBs and Packages

Cisco Crosswork Data Gateway has some pre-loaded MIBs and device packages. You can download them to obtain a tarball of the custom MIBs and device packages from the Crosswork UI, add more custom MIBs and device packages and re-upload them to the Crosswork.

System MIB Packages and System Device Packages are downloadable only. This is only if you want to know the abilities that already exists in the system. These cannot be modified.

If you have a new version, you can delete the existing one and upload the new one.

Follow the below steps to download custom software packages from Crosswork UI.

Step 1 From the main menu, choose **Admin** > **Data Gateway Global Settings**.

Step 2 From **Custom Software** pane, choose based on the following table:

If you want to download	Click...
Custom MIB Package	Download Custom MIB Package
System MIB Package	Download System MIB Package
System Device Package	Download System Device Package

Step 3 In the download window, navigate to the location where you want to download the file and click **Save**.

What to do next

To add new MIBs/Yangs, follow the steps:

1. Extract the package and add new files.

2. Run the package through the offline tool as explained at Use Custom MIBs and Yangs on Cisco DevNet to ensure that it can be uploaded to Crosswork.
3. Tar it back as custom-mib-packages.tar
4. Run XZ utility to compress it to custom-mib-packages.tar.xz
5. Upload the package back into Crosswork by following the steps described at Add a Custom Software Package, on page 184.

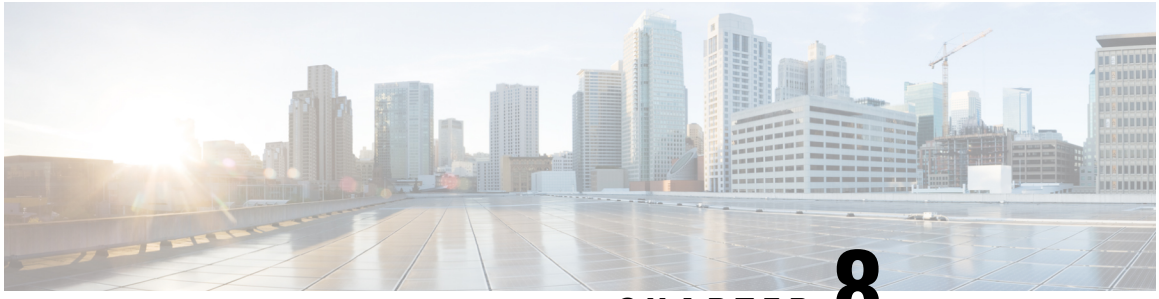

CHAPTER 8

Configure Collection

This section contains the following topics:

- Collection Service Overview, on page 189
- Collection Considerations, on page 189
- Prerequisites for Device Model Driven Telemetry, on page 190
- About Collection Jobs, on page 192
- Collection Jobs, on page 193
- Monitoring Collection Jobs, on page 201

Collection Service Overview

Networks maintain a large amount of data that spans thousands of devices. The Cisco Crosswork Optimization Engine collects and manages that flow of data via Cisco Crosswork Data Gateway in a multi-vendor environment to provide a real-time publish/subscribe model infrastructure to Cisco Crosswork Network Automation applications. The Collection Service is highly scalable in order to meet the performance demands of the Cisco Crosswork Optimization Engine applications.

Multiple applications requesting same data overload network devices causing outages. Cisco Crosswork Data Gateway's **Collection Optimization** feature tackles this problem by optimizing collection requests. Thereby, reducing redundant data collections.

Users with administrative privileges can monitor Collection Service status and performance, and start/stop/restart it or its underlying services, using the Cisco Crosswork Optimization Engine user interface. You can also collect logs and performance metrics for this service. For help with these tasks, see Manage Cisco Crosswork Network Automation, on page 125.

Collection Considerations

MDT Collection

When Cisco NSO is used in conjunction with Cisco Crosswork Optimization Engine, the telemetry configurations are pushed to the devices by Cisco NSO.

If you do not plan to use to use Cisco NSO, you must apply the telemetry configuration on your devices. See Prerequisites for Device Model Driven Telemetry, on page 190.

Note The default MDT collector port is 9010.

Device Limits

Cisco Data Gateway collection supports 1000 devices. If your network requires collection of more than 1000 devices, multiple Cisco Data Gateways must be deployed.

Prerequisites for Device Model Driven Telemetry

The Cisco Crosswork Optimization Engine Collection Service configures telemetry as needed on the devices enrolled within the service. Telemetry configuration must be done on PCCs or provider edge routers.

Note Cisco Crosswork Optimization Engine uses the Cisco-IOS-XR-infra-tc-oper YANG module for MDT collection.

Note If an operator configures telemetry directly on the same devices either manually or through some mechanism outside of the Collection Service, the commands must not contain the keyword `CW`. The keyword `CW` is reserved for use by the Collection Service. In particular, the following commands must not contain the keyword `CW` when configured outside of the Collection Service:

```
destination-group
sensor-group
subscription
        sensor-group-id
        destination-id
```

For example (invalid telemetry configuration):

```
telemetry model-driven
 destination-group CW_1b4ac245d863cf3e787d42bae97f1d18dd300d5e
```

For more information, see the telemetry configuration documentation for your particular device (for example: Telemetry Configuration Guide for Cisco ASR 9000)

The default MDT collector port is 9010. Cisco Data Gateway collection supports 1000 devices. If your network requires collection of more than 1000 devices, multiple Cisco Data Gateways must be deployed. See Collection Considerations, on page 189.

Valid Telemetry Configuration

The following sample output shows a *valid* telemetry configuration on a device when configured outside of the Collection Service. If using a single interface network, then the IP address is the management IP address. In a dual interface network, then the IP address should be the data IP address.

```
telemetry model-driven
destination-group OE_43dc8a5ea99529715899b4f5218408a785e40fce
vrf default
```

```
address-family ipv4 192.168.0.3 port 9010
encoding self-describing-gpb
protocol tcp
!
!
destination-group OE_4b3c69a200668b0a8dc155caff295645c684a8f8
vrf default
  address-family ipv4 192.168.0.3 port 9010
   encoding self-describing-gpb
   protocol tcp
  !
!
sensor-group OE_43dc8a5ea99529715899b4f5218408a785e40fce
  sensor-path Cisco-IOS-XR-infra-tc-oper:traffic-collector/afs/af/counters/tunnels/tunnel
!
sensor-group OE_4b3c69a200668b0a8dc155caff295645c684a8f8
  sensor-path
Cisco-IOS-XR-infra-tc-oper:traffic-collector/vrf-table/default-vrf/afs/af/counters/prefixes/prefix
!
subscription OE_43dc8a5ea99529715899b4f5218408a785e40fce
  sensor-group-id OE_43dc8a5ea99529715899b4f5218408a785e40fce sample-interval 60000
  destination-id OE_43dc8a5ea99529715899b4f5218408a785e40fce
!
subscription OE_4b3c69a200668b0a8dc155caff295645c684a8f8
  sensor-group-id OE_4b3c69a200668b0a8dc155caff295645c684a8f8 sample-interval 60000
  destination-id OE_4b3c69a200668b0a8dc155caff295645c684a8f8
!
```

Note The **`sample-interval`** can be changed depending on the size of your network. It is defined in milliseconds and determines how fast you want the data to be pushed out.

Confirm that all PCCs or provider edge routers have telemetry configured and report data to Crosswork Optimization Engine. For example, routers should report prefix and tunnel counters:

```
RP/0/RP0/CPU0:PE1#show traffic-collector ipv4 counters prefix
Thu Jul 11 08:32:32.993 UTC
Prefix Label Base rate TM rate State
(Bytes/sec) (Bytes/sec)
------------------------------ -------- ------------ ------------ ----------
192.168.0.1/32 16001 1 0 Active
192.168.0.2/32 16002 1 0 Active
192.168.0.3/32 16003 1 0 Active
192.168.0.4/32 16004 2 0 Active
192.168.0.6/32 16006 501023 501021 Active
192.168.0.7/32 16007 17320774 17320772 Active
192.168.0.8/32 16008 3737825 3737823 Active
192.168.0.9/32 16097 3 0 Active
192.168.0.10/32 16096 2 0 Active
```

```
RP/0/RP0/CPU0:PE1#show traffic-collector ipv4 counters tunnel
Thu Jul 11 08:32:20.746 UTC
Interface Base rate Base rate State
(Packet/sec) (Bytes/sec)
------------------------------ ------------ ------------ ----------
srte_c_102_ep_192.168.0.7 0 0 Active
```

Cisco IOS XR devices that are onboarded through telemetry must have the following configuration settings on the device to ensure that NETCONF and SSH work correctly:

```
ssh server v2
ssh server vrf default
ssh server netconf vrf default
ssh server rate-limit 600
ssh server session-limit 1024
netconf-yang agent ssh
```

Cisco IOS XR devices that are onboarded through SNMP must have SNMP enabled on the device. The following is an example of an SNMP configuration on a Cisco IOS XR device:

```
snmp-server community public RO
```

Please note that, currently, Cisco Crosswork Optimization Engine does not itself support execution of EXEC privilege commands, such as **enable**, on devices. These types of commands must be executed using the device console or other means.

About Collection Jobs

Note This section describes Crosswork Data Gateway features that are supported by various products in the Crosswork Network Automation suite. Please note that Cisco Crosswork Optimization Engine applications are, by default, automatically subscribed for required SNMP and MDT data collection. Cisco Crosswork Optimization Engine receives the data to internal Kafka and creates new collection jobs. Keep in mind that Crosswork Data Gateway extended collection features, MIBs, and other data described in this section may not apply to Cisco Crosswork Optimization Engine.

As mentioned earlier, Crosswork Data Gateway pulls functional images from the Crosswork. Each functional image represents a collection type. You can create multiple jobs for a given collection type. A collection job describes what task a Crosswork Data Gateway is expected to perform. Crosswork receives the data collection requests via these collection jobs and assigns to a Crosswork Data Gateway instance to serve the request.

You can collect more than one type of data at a time by using separate collection jobs.

For each collection job you create, Crosswork Data Gateway executes the collection request and deposits the collected data in the preferred data destination(s).

Crosswork Data Gateway lets you create three types of collection jobs:

CLI Collection Job

Enables CLI-based data collection (such as device configuration) from the network devices. The CLI collector uses XDE/PAL to collect device data for a given CLI. Only **show** commands are supported for this type of collection job.

SNMP Collection Job

Enables SNMP-based data collection based on the OIDs supported on the devices.

Supported SNMP versions include SNMPv1, SNMPv2c, and SNMPv3 for data polling and traps.

MDT Collection Job

Collects model driven telemetry data streamed from the device to the Crosswork Data Gateway.

Note

1. Crosswork Data Gateway drops incoming southbound traffic if there is no corresponding (listening) collection job request for the same. It also drops data/SNMP traps received from an unsolicitied device (i.e., not attached to Crosswork Data Gateway). Crosswork Data Gateway records this in log and notifies Crosswork.
2. Polled data cannot be requested from the device until Crosswork Data Gateway is ready to process and transmit the data. If it cannot keep up with the amount of data, it sends an error to northbound interface indicating when the throttling began and condition cleared.
3. The status of a collection job shows as **Unknown** if the job has not received any data. For example, collection jobs for traps will show **Unknown** when traps are not configured on the device or when there are no changes to trigger a traps notification for the device.

Collection Jobs

This section contains sample collection job payloads for the following collection profiles:

- CLI Collection Job, on page 193
- SNMP Collection Jobs, on page 194
- MDT Collection Job, on page 200

CLI Collection Job

Crosswork Data Gateway supports CLI-based data collection from the network devices. It uses XDE/PAL to collect device data for a given CLI. Only show commands are supported for this type of collection job.

Note

- The initial status for all the collection jobs in the UI is Unknown. Upon receiving a CLI collection job, Cisco Crosswork Data Gateway performs basic validations on it. If the collection job is valid, its status changes to Successful, else it changes to Failed.
- Device should not have any banner configuration for CLI collection to work properly. Please refer to device documentation on how to turn this off.
- The value of **Cadence** is in seconds. It should be set either to 0 to indicate the sensor configured to be collected only once.

  OR

  It should be >= 60 (i.e. at least 1 minute) up to 2764800 seconds (i.e. at most 32 days) max, indicating how frequently configured sensor data should be collected.
- When collection from a device is skipped due to previous execution still in progress, Cisco Crosswork Data Gateway raises a warning log. No alert is generated for this scenario.

Following is a CLI collection job sample:

```
{
  "collection_job": {
    "application_context": {
      "context_id": "collection-job1",
      "application_id": "APP1"
    },
    "collection_mode": {
      "lifetime_type": "APPLICATION_MANAGED",
      "collector_type": "CLI_COLLECTOR"
    },
    "job_device_set": {
      "device_set": {
        "devices": {
          "device_ids": [
            "658adb03-cc61-448d-972f-4fcec32cbfe8"
          ]
        }
      }
    },
    "sensor_input_configs": [
      {
        "sensor_data": {
          "cli_sensor": {
            "command": "show platform"
          }
        },
        "cadence_in_millisec": "tel:60000"
      }
    ],
    "sensor_output_configs": [
     {
        "sensor_data": {
          "cli_sensor": {
            "command": "show platform"
          }
        },
        "destination": {
          "destination_id": "1e71f2fb-ea65-4242-8efa-e33cec71b369",
          "context_id": "topic1"
        }
      }
    ]
  }
}
```

SNMP Collection Jobs

Crosswork Data Gateway supports SNMP-based data collection based on the OIDs supported on the devices.

The SNMP collector makes a poll request to Crosswork to get its configuration profile (a list of MIB objects to collect and a list of devices to fetch from). It determines the corresponding OIDs by looking up the pre-packaged list of MIB modules or the custom list of MIB modules.

Note MIBs are required only if the collection request references MIB TABLE names or SCALAR names. However, if the requests are OID-based, then MIBs are not required.

Once the OIDs are resolved, they are provided as input to the SNMP collectors.

The device packages can be imported into the Crosswork Data Gateway VM as described in Section Add a Custom Software Package, on page 184.

The following SNMP versions are supported:

- SNMPv1
- SNMPv2c
- SNMPv3

The table below lists supported privacy protocols and the value that needs to be given in the collection payload for SNMP and SNMP Trap collection jobs:

Protocol	SNMP Collection Payload	SNMP Trap Collection Payload
aes	AES	N/A
des56	DES	DES
3des	3DES	3DES
aes 128	AES128	AES128
aes 192	AES192 or CiscoAES192(Cisco specific)	AES192 or CiscoAES192(Cisco specific)
aes 256	AES256 or CiscoAES256(Cisco specific)	AES256 or CiscoAES256(Cisco specific)

Note

- The initial status for all the collection jobs in the UI is Unknown. Upon receiving a SNMP collection job, Cisco Crosswork Data Gateway performs basic validations on it. If the collection job is valid, its status changes to Successful, else it changes to Failed.
- The value of **Cadence** is in seconds. It should be set either to 0 to indicate the sensor configured to be collected only once.

  OR

  It should be >= 60 (i.e. at least 1 minute) up to 2764800 seconds (i.e. at most 32 days) max, indicating how frequently configured sensor data should be collected.
- When collection from a device is skipped due to previous execution still in progress, Crosswork Data Gateway raises a warning log. No alert is generated for this scenario.
- For SNMP v1/v2c, if the device details (such as host or community string) are incorrect in the payload, Crosswork Data Gateway ignores the traps received from the device and logs the a WARN message.

  In case of SNMP v3, if the device details (such as auth, priv, and security name details) are incorrect in the payload, Crosswork Data Gateway filters it out and hence, does not receive the trap. Thus, no WARN message is logged.

Sample Configurations on Device:

Version	Configuration
V1	snmp-server group group1 v1 snmp-server user user1 group1 v1 snmp-server host <host_ip> traps <community_string> udp-port 1062 For example, snmp-server host 172.29.194.78 traps test udp-port 1062 **Note** Version 1 is the default version used by the device.
V2c	snmp-server group group1 v2c snmp-server user user1 group1 v2c snmp-server host 172.29.194.142 traps version 2c v2test udp-port 1062
V3	snmp-server group group1 v3 auth notify user1 read user1 write user1 snmp-server view user1 1.3 included snmp-server user user1 group1 v3 auth md5 <password> priv aes 128 <password> snmp-server host 172.23.92.193 traps version 3 priv user1 udp-port 1062

The SNMP Collector supports the following operations:

- SCALAR
- TABLE
- MIB_WALK
- TRAP
- DEVICE_PACKAGE

These operations are defined in the sensor config (see payload sample below).

Note There is an optional **deviceParams** attribute **snmpRequestTimeoutMillis** (not shown in the sample payloads) that should be used if the device response time is very high. It's not recommended to use **snmpRequestTimeoutMillis** unless you are absolutely certain that your device response time is very high.

The value for snmpRequestTimeoutMillis should be specified in milliseconds:

Default value is 1500 milliseconds

Minimum value is 1500 milliseconds

However, there is no limitation on the maximum value of this attribute.

Following is an SNMP collection job sample:

```
{
  "collection_job": {
    "application_context": {
      "context_id": "collection-job1",
      "application_id": "APP1"
    },
    "collection_mode": {
      "lifetime_type": "APPLICATION_MANAGED",
      "collector_type": "SNMP_COLLECTOR"
    },
    "job_device_set": {
      "device_set": {
        "devices": {
          "device_ids": [
            "c70fc034-0cbd-443f-ad3d-a30d4319f937",
            "8627c130-9127-4ed7-ace5-93d3b4321d5e",
            "c0067069-c8f6-4183-9e67-1f2e9bf56f58"
          ]
        }
      }
    },
    "sensor_input_configs": [
      {
        "sensor_data": {
          "snmp_sensor": {
            "snmp_mib": {
              "oid": "1.3.6.1.2.1.1.3.0",
              "snmp_operation": "SCALAR"
            }
          }
        },
        "cadence_in_millisec": "60000"
      },
      {
        "sensor_data": {
          "snmp_sensor": {
            "snmp_mib": {
              "oid": "1.3.6.1.2.1.31.1.1",
              "snmp_operation": "TABLE"
            }
          }
        },
        "cadence_in_millisec": "60000"
      }
    ],
    "sensor_output_configs": [
      {
        "sensor_data": {
          "snmp_sensor": {
            "snmp_mib": {
              "oid": "1.3.6.1.2.1.1.3.0",
              "snmp_operation": "SCALAR"
            }
          }
        },
        "destination": {
          "destination_id": "4c2ab662-2670-4b3c-b7d3-b94acba98c56",
          "context_id": "topic1_461cb8aa-a16a-44b8-b79f-c3daf3ea925f"
        }
      },
      {
        "sensor_data": {
          "snmp_sensor": {
            "snmp_mib": {
```

```
                "oid": "1.3.6.1.2.1.31.1.1",
                "snmp_operation": "TABLE"
              }
            }
          },
          "destination": {
            "destination_id": "4c2ab662-2670-4b3c-b7d3-b94acba98c56",
            "context_id": "topic2_e7ed6300-fc8c-47ee-8445-70e543057f8a"
          }
        }
      ]
    }
  }
```

SNMP Traps Collection Job

SNMP traps are handled in a similar manner. Trap listeners listen on a port and then dispatch data to recipients (based on their topic of interest).

Note

- Device should have been pre-configured by the traps.
- Crosswork Data Gateway listens on UDP port 1062 for Traps.
- If the collection job is invalid, there is missing configuration on the device, or no trap is received, the status of the job remains "Unknown".

On receiving a trap, Crosswork Data Gateway does the following validations:

1. Check if any collection job is created for the device.
2. Checks the trap version and community string.
3. For SNMP v3, validates for user auth and priv protocol and credentials.

Following is an SNMP-Trap collection job sample:

```
{
  "collection_job": {
    "application_context": {
      "context_id": "collection-job1",
      "application_id": "APP1"
    },
    "collection_mode": {
      "lifetime_type": "APPLICATION_MANAGED",
      "collector_type": "TRAP_COLLECTOR"
    },
    "job_device_set": {
      "device_set": {
        "devices": {
          "device_ids": [
            "a9b8f43d-130b-4866-a26a-4d0f9e07562a",
            "8c4431a0-f21d-452d-95a8-84323a19e0d6",
            "eaab2647-2351-40ae-bf94-6e4a3d79af3a"
          ]
        }
      }
    },
    "sensor_input_configs": [
      {
```

```
        "sensor_data": {
          "trap_sensor": {
            "path": "1.3.6.1.6.3.1.1.4"
          }
        },
        "cadence_in_millisec": "60000"
      }
    ],
    "sensor_output_configs": [
      {
        "sensor_data": {
          "trap_sensor": {
            "path": "1.3.6.1.6.3.1.1.4"
          }
        },
        "destination": {
          "destination_id": "4c2ab662-2670-4b3c-b7d3-b94acba98c56",
          "context_id": "topic1_696600ae-80ee-4a02-96cb-3a01a2415324"
        }
      }
    ]
  }
}
```


Enabling Traps forwarding to external applications

As per the current implementation, in case of an SNMP Trap collection job, all traps are sent to the specified data destination even if the SNMP Trap OID is provided in the sensor path.

Therefore, it is recommended to have a single SNMP Trap collection job per device (with any OID as sensor path) as it would be enough to get all traps from that device.

To identify the type of trap from the data received on the destination, look for *oid* (OBJECT_IDENTIFIER, for example, `1.3.6.1.6.3.1.1.4.1.0`) and *strValue* associated to the *oid* in the OidRecords (application can match the OID of interest to determine the kind of trap).

Below are some sample values and a sample payload:

- Link up

  `1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.6.3.1.1.5.4`

- Link Down

  `1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.6.3.1.1.5.3`

- Syslog

  `1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.4.1.9.9.41.2.0.1`

- Cold Start

  `1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.6.3.1.1.5.1`


```
{
  "nodeIdStr": "BF5-XRV9K1.tr3.es",
  "nodeIdUuid": "C9tZ5lJoSJKf5OZ67+U5JQ==",
  "collectionId": "133",
  "collectionStartTime": "1580931985267",
  "msgTimestamp": "1580931985267",
  "dataGpbkv": [
    {
      "timestamp": "1580931985267",
```

```
      "name": "trapsensor.path",
      "snmpTrap": {
        "version": "V2c",
        "pduType": "TRAP",
        "v2v3Data": {
          "agentAddress": "172.70.39.227",
          "oidRecords": [
            {
              "oid": "1.3.6.1.2.1.1.3.0",
              "strValue": "7 days, 2:15:17.02"
            },
            {
              "oid": "1.3.6.1.6.3.1.1.4.1.0",  // This oid is the Object Identifier.
              "strValue": "1.3.6.1.6.3.1.1.5.3" // This is the value that determines the
kind of trap.
            },
            {
              "oid": "1.3.6.1.2.1.2.2.1.1.8",
              "strValue": "8"
            },
            {
              "oid": "1.3.6.1.2.1.2.2.1.2.8",
              "strValue": "GigabitEthernet0/0/0/2"
            },
            {
              "oid": "1.3.6.1.2.1.2.2.1.3.8",
              "strValue": "6"
            },
            {
              "oid": "1.3.6.1.4.1.9.9.276.1.1.2.1.3.8",
              "strValue": "down"
            }
          ]
        }
      }
    }
  ],
  "collectionEndTime": "1580931985267",
  "collectorUuid": "YmNjZjEzMTktZjFlOS00NTE5LWI4OTgtY2Y1ZmQxZDFjNWExOlRSQVBfQ09MTEVDVE9S",
  "status": {
    "status": "SUCCESS"
  },
  "modelData": {},
  "sensorData": {
    "trapSensor": {
      "path": "1.3.6.1.6.3.1.1.5.4"
    }
  },
  "applicationContexts": [
    {
      "applicationId": "APP1",
      "contextId": "collection-job-snmp-traps"
    }
  ]
}
```


MDT Collection Job

Crosswork Data Gateway supports data collection from network devices using Model-driven Telemetry (MDT) to consume telemetry streams directly from devices.

Note

- MDT collector retains the collection ID that comes as part of the telemetry proto for the device. This behavior is different from CLI and SNMP collectors which compute the collection ID based on the sequence number of the collection.
- MDT collection jobs require some configuration to be done on the device. This configuration is automatically taken care of by NSO.
- If there is some change (delete/update) in existing MDT jobs between backup and restore operations, Crosswork does not replay the jobs for config update on the devices as it involves Provider(NSO). You have to restore configs on provider/devices. Crosswork will just restore the jobs in database.

Monitoring Collection Jobs

Once a device is mapped to a Cisco Crosswork Data Gateway instance, the status of all the associated collection jobs is set to 'Accepted'.

From the **Collection Jobs** view, you can monitor the status of the collection jobs currently active on all the Cisco Crosswork Data Gateway instances enrolled with Cisco Crosswork Optimization Engine, such as system jobs and API-defined collection jobs.

From the navigation bar, choose **Admin** > **Collection Jobs**.

Figure 62: Add Optima image here 442649.jpg

Item	Description
Data Gateway Jobs Pane	Shows the list of all active collection jobs along with their status and application ID.
Job Details Pane	Shows the details of a particular job selected in the **Data Gateway Jobs** pane.

To view details of a collection job, select the collection job from the **Data Gateway Jobs** pane. The details of the selected job are displayed in the **Job Details** pane right next to the **Data Gateway Jobs** pane.

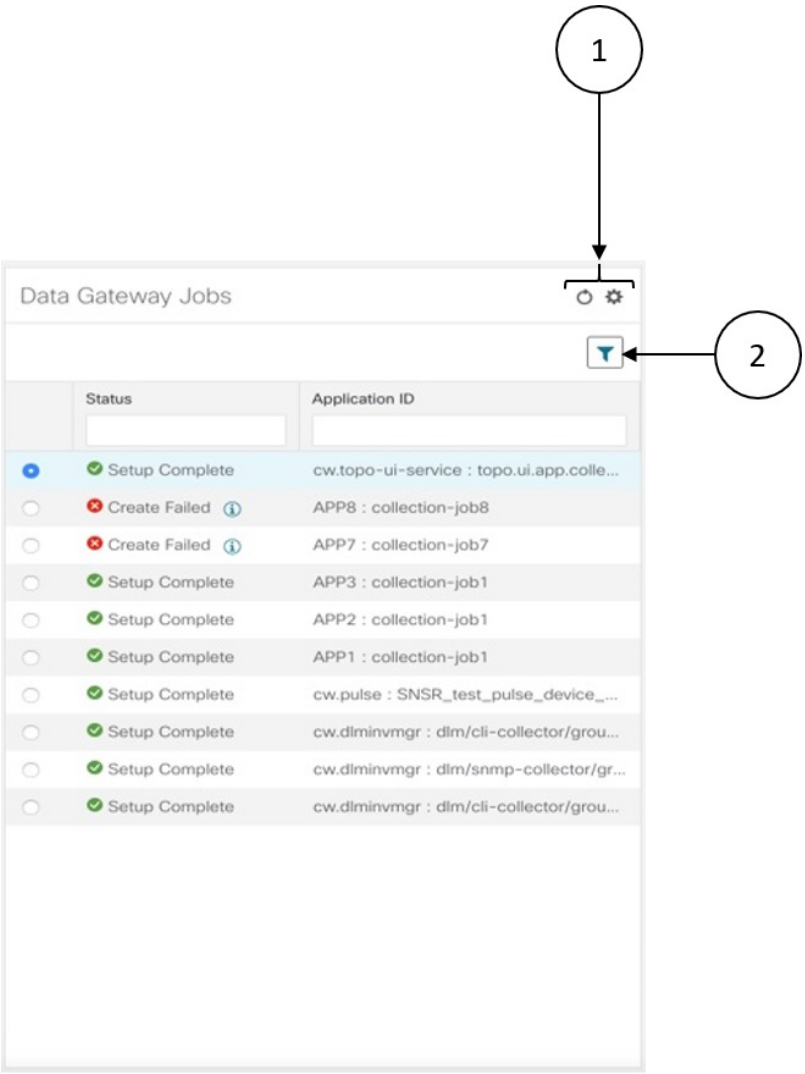

Item	Description
1	Click to refresh the **Data Gateway Jobs** window.
	Click to choose the columns to make visible in the **Data Gateway Jobs** window (see Set, Sort and Filter Table Data, on page 7).
2	Click to set filter criteria on one or more columns in the **Data Destinations** window.
	Click the **Clear All Filters** link to clear any filter criteria you may have set.

Data Gateway Jobs pane displays only the status and application ID.

Note

- `Create Failed` error means out of N devices , some devices detup failed. However, the collection would happen on the other devices. You can identify the device(s) causing this error by using `Control Status` API.
- If job creation failed on a particular device because of NSO errors, after fixing NSO errors , you have to manually change the administration state of the device first to "Down" and then "Up". However, doing so resets the collection on the device.

  Note

  Create/Delete failed errors are shown in a different screen pop up. Click to the job status to see details of the error.

- You may also try recreating the job using PUT collection job API with the same payload.

However, when you select a job, more details are displayed in the **Job Details** pane:

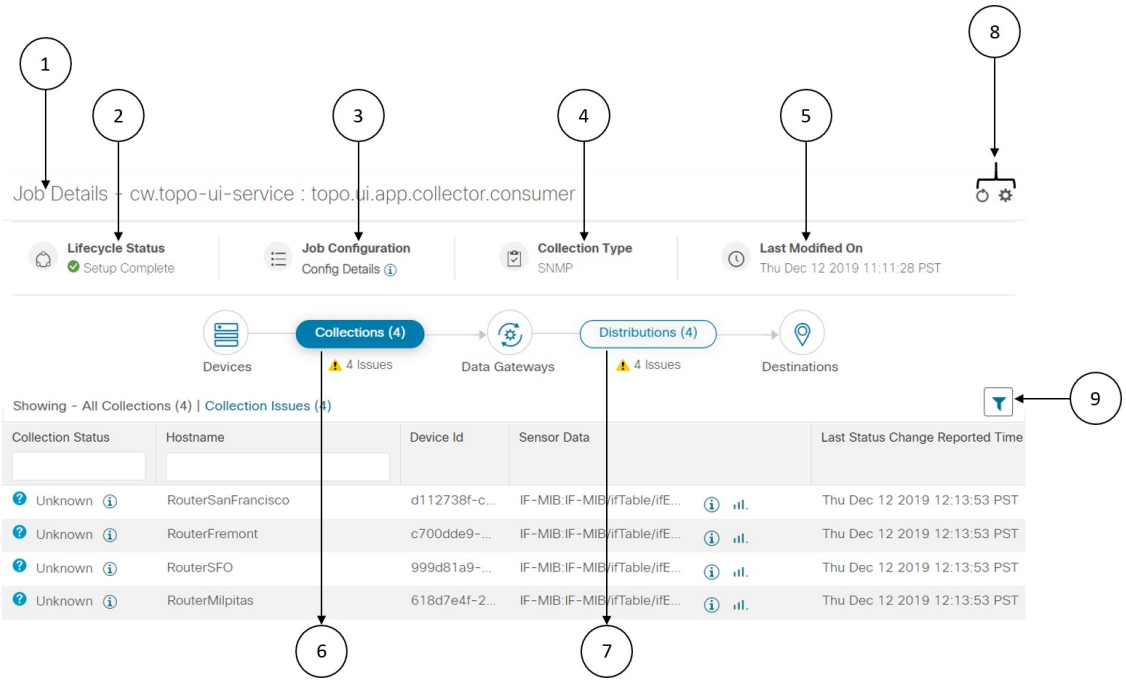

Item	Description
1	Application name and context associated with the collection job.
2	Lifecycle status of the collection job.

Item	Description
3	Job payload of the collection job that you pass in the REST API request. Click icon next to **Config Details** to view the job configuration. Crosswork Data Gateway lets you view configuration in two modes: • View Mode • Text Mode
4	Collection Type
5	Time and date of last modification of the collection job.
6	Collections (x): x refers to requested input collections that span device by sensor paths. The corresponding **(y) Issues** is the count of input collections that are in UNKNOWN or FAILED state.
7	Distributions (x): x refers to requested output collections that span device by sensor paths. The corresponding **(y) Issues** is the count of output collections that are in UNKNOWN or FAILED state.
8	Click to refresh the **Job Details** window.
	Click to choose the columns to make visible in the **Job Details** window (see Set, Sort and Filter Table Data, on page 7).
9	Click to set filter criteria on one or more columns in the **Job Details** window.
	Click the **Clear All Filters** link to clear any filter criteria you may have set.

Job Details pane displays the following details about a collection job:

Field	Description
Collection/Distribution Status	Status of the collection/distribution. It is reported on a on change basis from Cisco Crosswork Data Gateway. Click next to the collection/distribution status for details.
Hostname	Application with which the collection job is associated.
Device Id	Unique identifier of the device from which data is being collected.

Field	Description
Sensor Data	Sensor path Click to see collection/distribution summary. From the sensor data summary pop up you can copy the sensor data by clicking **Copy to Clipboard**. and Click to see collection/distribution metrics summary. The metrics are reported on cadence-basis i.e., once every 10 minutes by default. It shows the following metrics for a collection: • last_collection_time_msec • total_collection_message_count • last_device_latency_msec • last_collection_cadence_msec It shows the following metrics for a collection: • total_output_message_count • last_destination_latency_msec • last_output_cadence_msec • last_output_time_msec • total_output_bytes_count
Destination	Data destination for the job.
Last Status Change Reported Time	Time and date on which last status change was reported for that device sensor pair from Cisco Crosswork Data Gateway.

APPENDIX A

Configure Cisco Crosswork Data Gateway Base VM

This appendix describes how to configure a Cisco Crosswork Data Gateway Base VM.

This section contains the following topics:

- About Cisco Crosswork Data Gateway Base VM, on page 207
- Basic Concepts, on page 210
- Manage Users, on page 212
- View Current System Settings, on page 214
- Change Current System Settings, on page 215
- Monitor Cisco Crosswork Data Gateway Health, on page 222
- Troubleshooting, on page 228

About Cisco Crosswork Data Gateway Base VM

A Cisco Crosswork Data Gateway instance is created as a standalone VM and can be geographically separate from the controller application (the controller application could be Crosswork Cloud or a Crosswork On-Prem application, such as Cisco Crosswork Optimization Engine). This Base VM is capable of connecting to the controller application and enable data collection from the network.

Crosswork orchestrates the collection from the distributed Cisco Crosswork Data Gateway VM instances.

The Cisco Crosswork Data Gateway VM is delivered as an OVA file and the additional functional images are delivered as Docker images.

Base VM Contents

The Base VM (OVA) is pre-packaged with basic functionality required to reach the controller application.

The Cisco Crosswork Data Gateway VM (OVA) contains the following pre-packaged contents:

- Cisco hardened Ubuntu distribution of Linux
- Cisco Crosswork Data Gateway services:
  - Vitals Monitor - Monitors resource usage on the VM.

- Controller Gateway – Establishes trusted connection with the controller application via the Controller Gateway and downloads functional images and configuration files.
- Image Manager – Coordinates between the Cisco Crosswork Data Gateway and the controller application to download functional images and configuration files.
- Route Manager – Directs traffic to devices on different south-bound destinations and also connects to the controller application and data devices via the north-bound interface.
- Docker IPv6nat - Programs IPv6 routes for docker containers.

Note Functional images (CLI, SNMP, and MDT collectors) are not included in the Base VM. They are downloaded by Cisco Crosswork Data Gateway from the controller application after successful authentication and bootstrap.

Log In and Log Out

You can use either of the following two ways to access Cisco Crosswork Data Gateway:

- Access Cisco Crosswork Data Gateway Through vCenter, on page 208
- Access Cisco Crosswork Data Gateway Via SSH, on page 209

Access Cisco Crosswork Data Gateway Through vCenter

Follow these steps to log in via vCenter:

Step 1 Locate the VM in vCenter and then right click and select **Open Console**.

The Cisco Crosswork Data Gateway flash screen comes up.

Step 2 Enter username (`dg-admin` or `dg-oper` as per the role assigned to you) and the corresponding password (the one that you created during installation process) and press **Enter**.

```
Cisco Crosswork Data Gateway

  #####  ######  #######  #####   #####  #     # ####### ######  #    #
 #     # #     # #     # #     # #     # #  #  # #     # #     # #   #
 #       #     # #     # #       #       #  #  # #     # #     # #  #
 #       ######  #     #  #####   #####  #  #  # #     # ######  ###
 #       #   #   #     #       #       # #  #  # #     # #   #   #  #
 #     # #    #  #     # #     # #     # #  #  # #     # #    #  #   #
  #####  #     # #######  #####   #####   ## ##  ####### #     # #    #

Copyright (c) 2019 by Cisco Systems, Inc.
Version: 1.1.0 (branch dg110dev - build number 245)
Built on: Nov-20-2019 00:06 AM UTC

Password:
```

Access Cisco Crosswork Data Gateway Via SSH

Note The SSH process is protected from brute force attacks by blocking the client IP after a number of login failures. Failures such as incorrect username or password, connection disconnect, or algorithm mismatch are counted against the IP. Up to 4 failures within a 20 minute window will cause the client IP to be blocked for at least 7 minutes. Continuing to accumulate failures will cause the blocked time to be increased. Each client IP is tracked separately.

Follow these steps to login via SSH.

Step 1 Run the following command:

ssh <username>@<ManagementNetworkIP>

where **ManagementNetworkIP** is the management network IP address.

For example,

To login as adminstrator user: **ssh dg-admin@<ManagementNetworkIP>**

To login as operator user: **ssh dg-oper@<ManagementNetworkIP>**

The following Cisco Crosswork Data Gateway flash screen opens prompting for password:

```
Cisco Crosswork Data Gateway

  #####  ######  #######  #####   #####  #     # ####### ######  #    #
 #     # #     # #     # #     # #     # #  #  # #     # #     # #   #
 #       #     # #     # #       #       #  #  # #     # #     # #  #
 #       ######  #     #  #####   #####  #  #  # #     # ######  ###
 #       #   #   #     #       #       # #  #  # #     # #   #   #  #
 #     # #    #  #     # #     # #     # #  #  # #     # #    #  #   #
  #####  #     # #######  #####   #####   ## ##  ####### #     # #    #

Copyright (c) 2019 by Cisco Systems, Inc.
Version: 1.1.0 (branch dg110dev - build number 245)
Built on: Nov-20-2019 00:06 AM UTC

Password:
```

Step 2 Input the corresponding password (the one that you created during installation process) and press **Enter**.

Use the Interactive Console

Cisco Crosswork Data Gateway launches an interactive console upon successful login. The interactive console displays the Main Menu as shown in the following figure:

Note The Main Menu shown here corresponds to **dg-admin** user. It is different for **dg-oper** user as the operator does not have same privileges as the adminstrator. See Supported User Roles, on page 212.

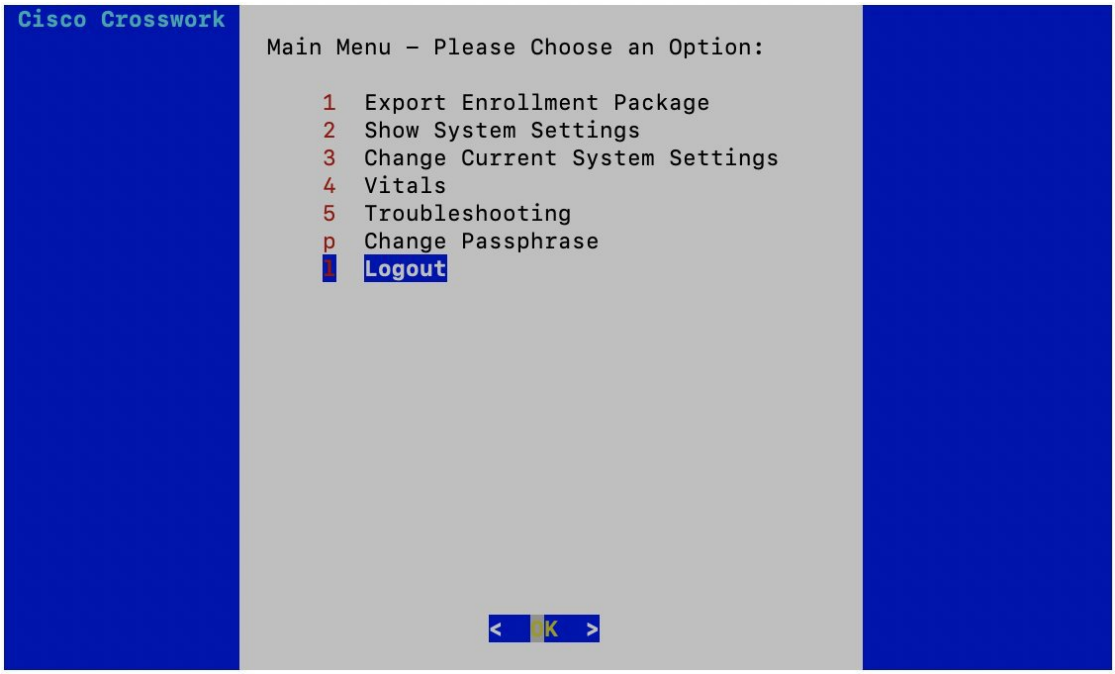

The Main Menu presents the following options:

1. Export Enrollment Package
2. Show System Settings
3. Change Current System Settings
4. Vitals
5. Troubleshooting

p. Change Passphrase

l. Logout

Basic Concepts

Cisco Crosswork Data Gateway makes extensive use of certain concepts. It is helpful to be familiar with them before you get started.

Cisco Crosswork Data Gateway Components

Cisco Crosswork Data Gateway has the following five main components or services:

- Controller Gateway, on page 211
- Image Manager, on page 211
- Vitals Monitor, on page 211

- Route Manager, on page 212
- Docker IPv6nat, on page 212

Controller Gateway

Controller Gateway is the component responsible for all the interaction between a Cisco Crosswork Data Gateway instance and its controller application. It manages the session creation with the controller application and makes sure all the payloads and responses are signed and verified for integrity. Components such as Image Manager, Vitals Monitor, and Route manager interact via Controller Gateway with the controller application to exchange the details those components need.

Note When the Controller Gateway stops, any alerts are not updated in cdg-alerts.log. However, when it starts, it sends an alert that it has started. This is because all the alerts go through the Controller Gateway and if it is down, the controller application won't receive the alerts. To access log files, see Run show-tech, on page 232.

Image Manager

The Image Manager starts up when Cisco Crosswork Data Gateway VM boots. It downloads the functional images from the repository as instructed by the controller application and brings up the services.

It has the following responsibilities:

- Periodically pull boot-config file from the controller application via Controller Gateway.
- Based on the boot-config and local images metadata cache, determine if the functional images and docker-compose file need to be downloaded.
- Send appropriate alerts to the controller application, if there are issues while processing the boot-config.
- Stop and remove any services that are no longer called for in the latest boot-config.
- Cleanup the local images metadata cache to keep it synchronized with the latest boot-config received from the controller application.
- Downloads collectors environment and other files that facilitate establishment of connection between collectors and Crosswork.
- Downloads system device packages and MIB packages required by the collectors from Crosswork.
- Downloads custom software to the collectors when uploaded via Crosswork UI.

Note Functional images are downloaded only when there is a change in boot-config response.

Vitals Monitor

The Vitals Monitor monitors the health and vitals of the Cisco Crosswork Data Gateway VM. It collects the CPU, memory, disk usage, docker containers metrics, etc. and aggregates this information in a file on the host filesystem.

For more information, see Monitor Cisco Crosswork Data Gateway Health, on page 222.

Route Manager

Route Manager manages south-bound routes to devices and north-bound routes to data destinations based on add/delete requests from collector upon the updates of inventory and collection jobs.

Route manager adds/deletes the static routes by comparing the existing routes configured on the VM with the routes configuration. This configuration is pushed to the Route Manager by the controller application in case of Crosswork On-Premise deployment.

Appropriate alerts are sent to the controller application if there is any failure in processing route request.

Docker IPv6nat

docker-ipv6nat is a special process that programs ipv6 routes for docker containers.

Manage Users

This section contains the following topics:

- Supported User Roles, on page 212
- Change Password, on page 214

Supported User Roles

Cisco Crosswork Data Gateway supports only two users with the following user roles:

- **Administrator**: One default user with administrator role is created when Cisco Crosswork Data Gateway is brought up for the first time. This user cannot be deleted and has both read and write privileges such as start/shut down Cisco Crosswork Data Gateway, register an application, apply authentication certificates, configure server settings, and perform kernel upgrade.
- **Operator**: This user is also created by default during the initial VM bring up. Operator can review the state/health of theCisco Crosswork Data Gateway, retrieve health/error logs, receive error notifications and run connectivity tests between Cisco Crosswork Data Gateway instance and the output destination.

Note

- Both users' credentials are configured during Cisco Crosswork Data Gateway installation.
- Users are locally authenticated.

The following table shows the permissions available to each role:

Table 17: Permissions Per Role

Permissions	Administrator	Operator
Export enrollment package	✓	✓
Show system settings		

Permissions	Administrator	Operator
• Management and South/North-bound Data Addresses • NTP • DNS • Proxy • UUID • Syslog • Certificates	✓	✓
Change Current System Settings		
• Configure NTP • Configure DNS • Configure Control Proxy • Configure Static Routes • Configure Syslog • Create new SSH keys • Import Certificate	✓	✕
Vitals		
• Docker Containers • Docker Images • Controller Reachability • NTP Reachability • Route Table • ARP Table • Network Connections • Disk Space Usage	✓	✓
Troubleshooting		
Ping a Host	✓	✓
Traceroute to a Host	✓	✓
NTP Status	✓	✓

Permissions	Administrator	Operator
System Uptime	✓	✓
Run show-tech	✓	✓
Remove All Collectors and Reboot VM	✓	×
Reboot VM	✓	×
Change Passphrase	✓	✓

Change Password

Both Adminstrator and Operator users can change their own passphrases but not each others'.

Follow these steps to change your passphrase:

Step 1 From the Main Menu, select **p Change Passphrase** and click **OK**.

Step 2 Input your current password and press Enter.

```
Changing password for dg-admin.
(current) UNIX password:
```

Step 3 Enter new password and press Enter. Re-type the new password and press Enter.

```
Changing password for dg-admin.
(current) UNIX password:
Enter new UNIX password:
Retype new UNIX password:
```

View Current System Settings

Cisco Crosswork Data Gateway allows you to view the following settings:

Follow these steps to view the current system settings:

Step 1 From the Main Menu, select **2 Show System Settings**, as shown in the following figure:

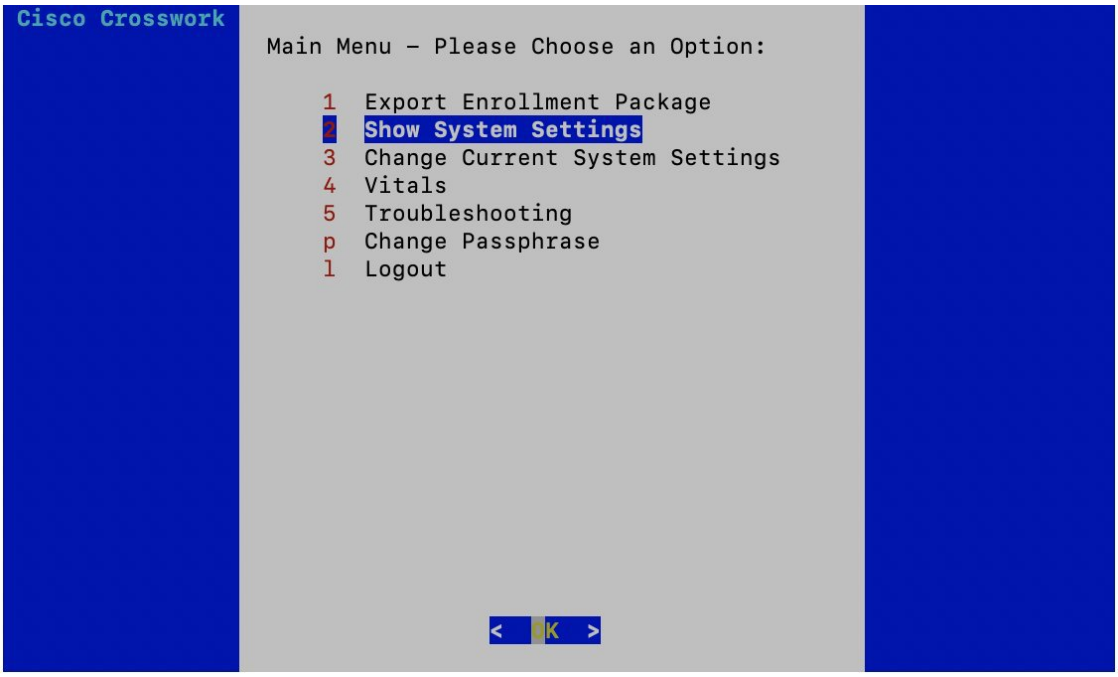

Step 2 Click **OK**. The **Show Current System Settings** menu opens.

Step 3 Select the setting you want to view.

Step 4 Click **OK**. Cisco Crosswork Data Gateway displays the selected setting.

After you are done viewing the settings, press any key to return to the **Show Current System Settings** menu.

To return to the Main Menu, select **x Exit Menu** and click **OK**.

Change Current System Settings

Note

- Cisco Crosswork Data Gateway System settings can only be configured by the Administrator.

Cisco Crosswork Data Gateway allows you to change the following settings:

- NTP
- DNS
- Control Proxy
- Static routes
- Syslog
- SSH keys

- Certificate

Follow these steps to change the current system settings:

Step 1 From the Main Menu, select **3 Change Current System Settings**, as shown in the following figure.

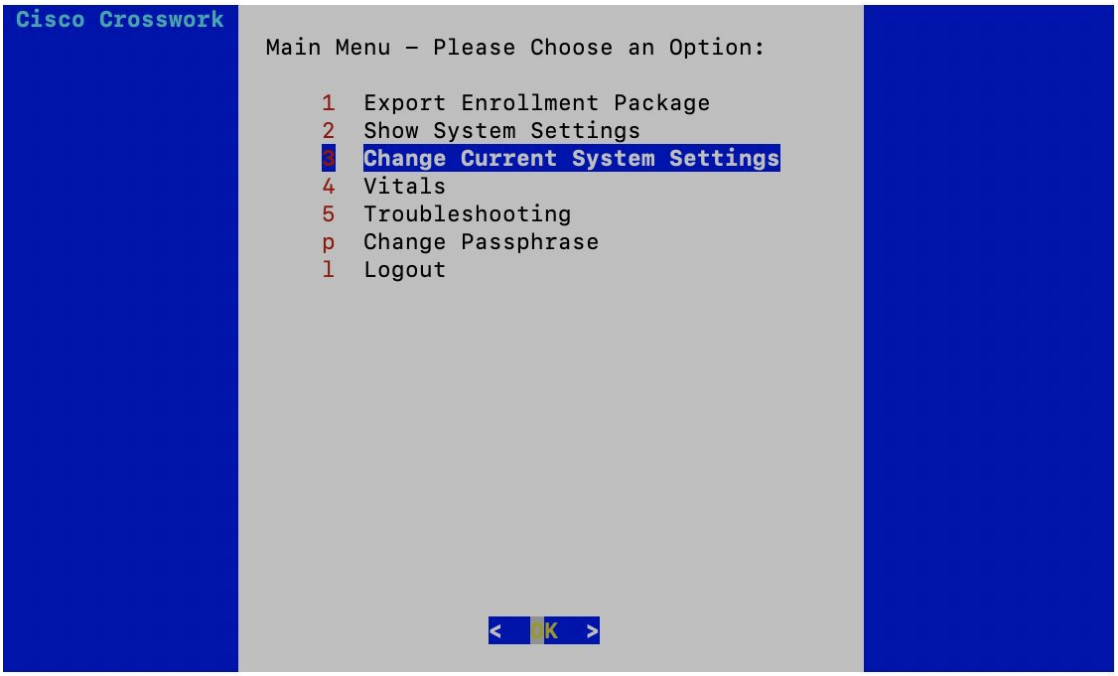

Step 2 Click **OK**. The **Change System Settings** menu opens.

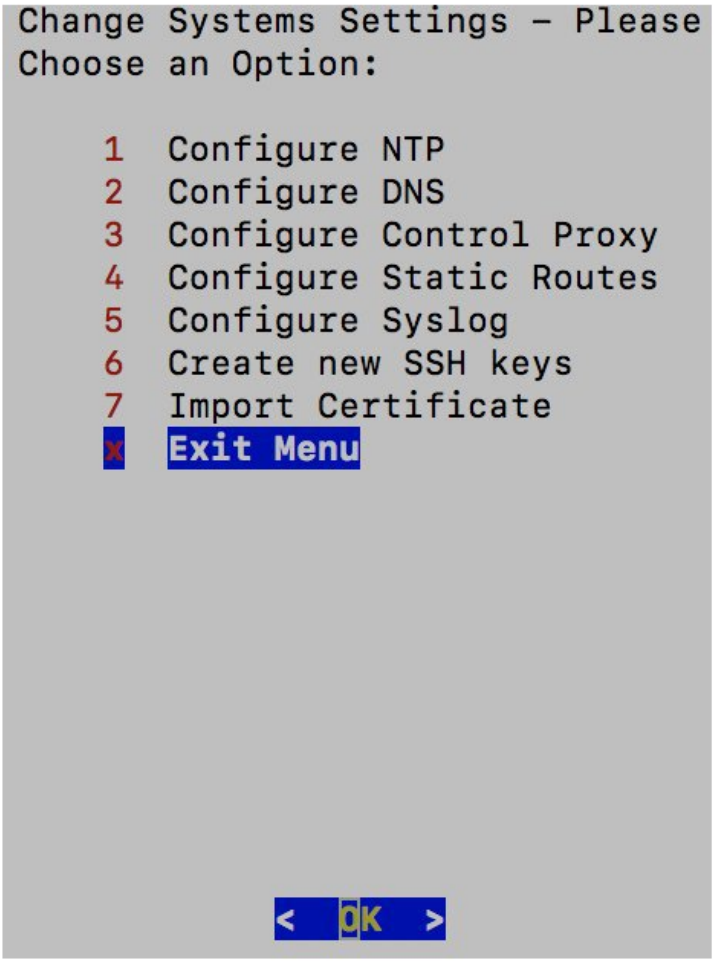

Step 3 Select the setting you want to change.

Step 4 Click **OK**. Cisco Crosswork Data Gateway prompts you to input new value for the selected setting.

Step 5 After you have entered the new settings, click **OK** to save the settings and return to the **Change System System Settings** menu.

To return to the Main Menu, select **x Exit Menu** and click **OK**.

Configure NTP

Step 1 From the **Change Current System Settings** Menu, select **1 Configure NTP** and click **OK**.

Step 2 Enter the new NTP server.

Step 3 Click **OK** to save the settings.

Configure DNS

Step 1 From the **Change Current System Settings** menu, select **2 Configure DNS** and click **OK**.

Step 2 Enter the new DNS domain and server address.

Step 3 Click **OK** to save the settings.

Configure Control Proxy

Step 1 From the **Change Current System Settings** menu, select **3 Configure Control Proxy** and click **OK**.

Step 2 Enter the new Proxy server URL and the exception list.

Step 3 Click **OK** to save the settings.

Configure Static Routes

In Cisco Crosswork Data Gateway, the static routes are configured when the Route Manager receives add/delete requests from the collectors. The **Configure Static Routes** option from the main menu can be used for troubleshooting purpose.

Note Static routes configured using this option are lost when the Cisco Crosswork Data Gateway reboots.

Add Static Routes

Step 1 From the **Change Current System Settings** menu, select **4 Configure Static Routes** and click **OK**.

Step 2 To add a static route, select **a Add** and click **OK**.

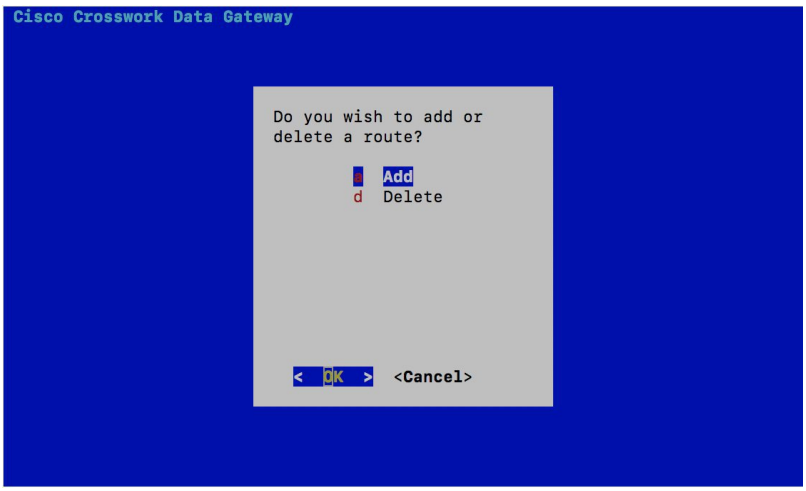

Step 3 Select the interface for which you want to add a static route and click **OK**.

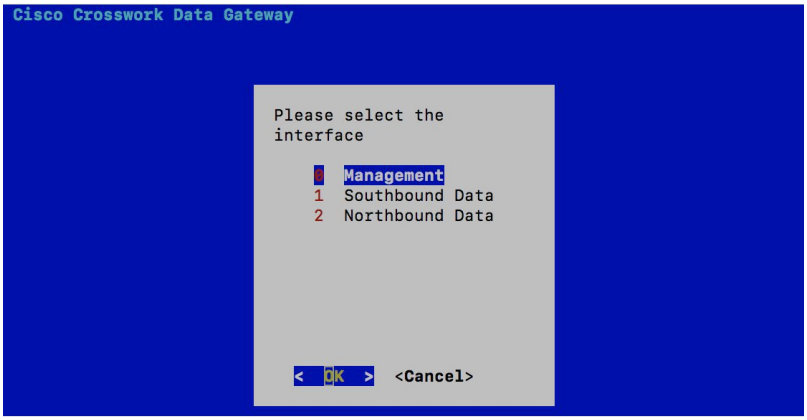

Step 4 Select the IP address version for which you want to add a route and click **OK**.

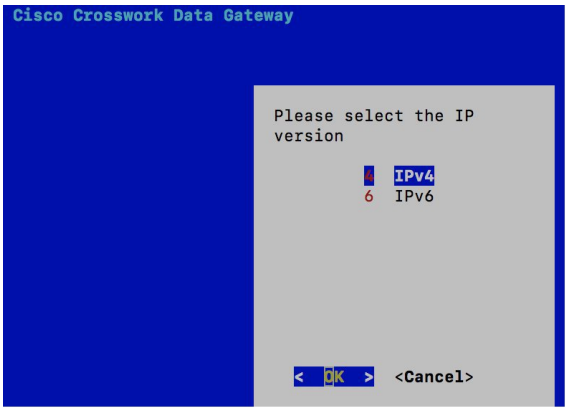

Step 5 Enter IPv4/IPv6 subnet in CIDR format when prompted.

Step 6 Click **OK** to save the settings.

Delete Static Routes

Step 1 From the **Change Current System Settings** Menu, select **4 Configure Static Routes** and click **OK**.

Step 2 To delete a static route, select **d Delete** and click **OK**.

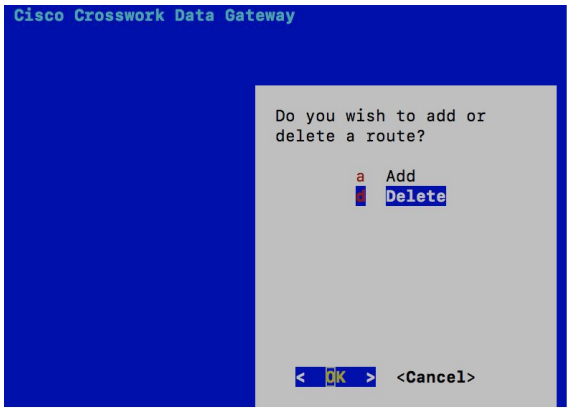

Step 3 Select the interface for which you want to delete a static route and click **OK**.

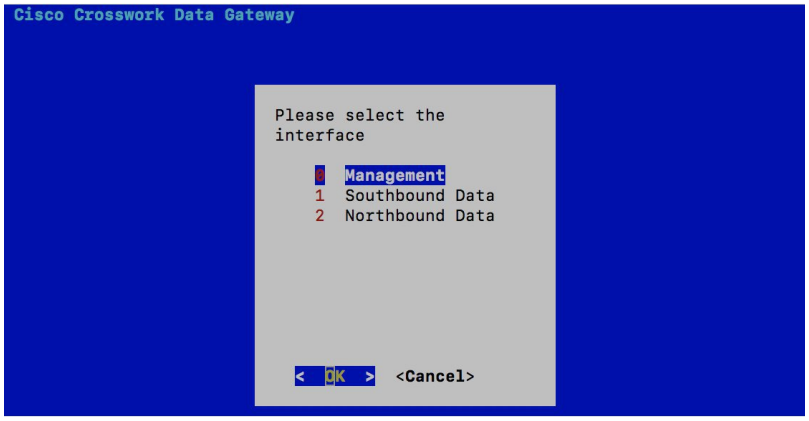

Step 4 Select the IP address version for which you want to delete a route and click **OK**.

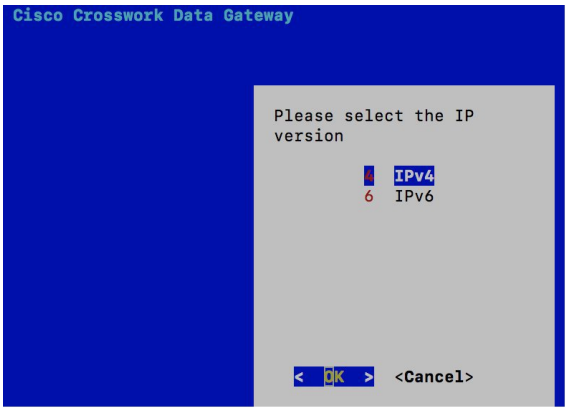

Step 5 Enter IPv4/IPv6 subnet in CIDR format.

Step 6 Click **OK** to save the settings.

Configure Syslog

Note For any Syslog server configuration with IPv4/IPv6 support for different linux distributions, please refer your system administrator and configuration guides.

Step 1 From the **Change Current System Settings** Menu, select **5 Configure Syslog** and click **OK**.

Step 2 Enter the new values for the following syslog attributes:.

- Server address: IPv4 or IPv6 address of a syslog server accessible from the management interface. If you are using an IPv6 addres, it must be surrounded by square brackets ([1::1]).
- Port: Port number of the syslog server
- Protocol: Use UDP, TCP, or RELP when sending syslog.
- Use Syslog over TLS?: Use TLS to encrypt syslog traffic.
- TLS Peer Name: Syslog server's hostname exactly as entered in the server certificate SubjectAltName or subject common name.
- Syslog Root Certificate File URI: PEM formatted root cert of syslog server retrieved using SCP.
- Syslog Certificate File Passphrase: Password of SCP user to retrieve Syslog certificate chain.

Step 3 Click **OK** to save the settings.

Create New SSH Keys

Step 1 From the **Change Current System Settings** Menu, select **6 Create new SSH keys**.

Step 2 Click **OK**. Crosswork Data Gateway launches an auto-configuration process that generates new SSH keys.

Import Certificate

Updating any certificate other than Controller Signing Certificate causes a collector restart.

Step 1 From the **Change Current System Settings** Menu, select **7 Import Certificate** and click **OK**.

Step 2 Select the certificate you want to import and click **OK**.

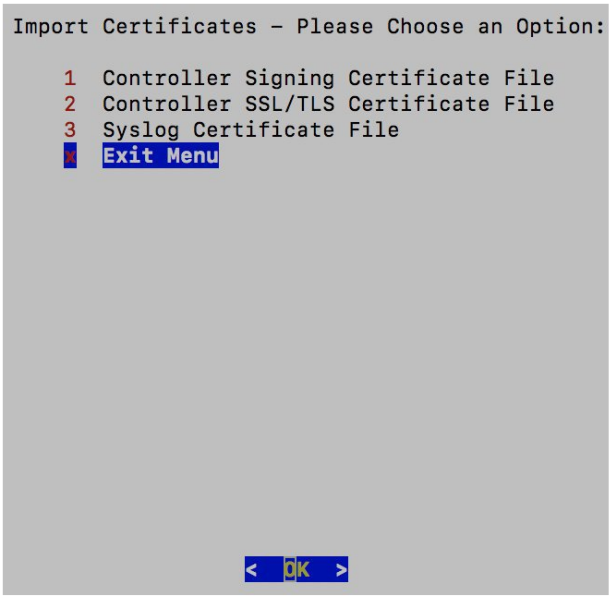

Step 3 Enter SCP URI for the selected certificate file and click **OK**.

Step 4 Enter passphrase for the SCP URI and click **OK**.

Monitor Cisco Crosswork Data Gateway Health

This section contains the following topics:

- Vitals Monitor, on page 223
- View Cisco Crosswork Data Gateway Vitals, on page 223
- collector-vitals Service, on page 226

Vitals Monitor

The Vitals Monitor component of Cisco Crosswork Data Gateway enables you to view vitals for the following:

1. Docker containers
2. Docker images
3. Controller reachability
4. NTP reachability
5. Route table
6. ARP table
7. Network connections
8. Disk space usage

View Cisco Crosswork Data Gateway Vitals

Follow these steps to view Cisco Crosswork Data Gateway vitals:

Step 1 From the Main Menu, select **4 Vitals** and click **OK**.

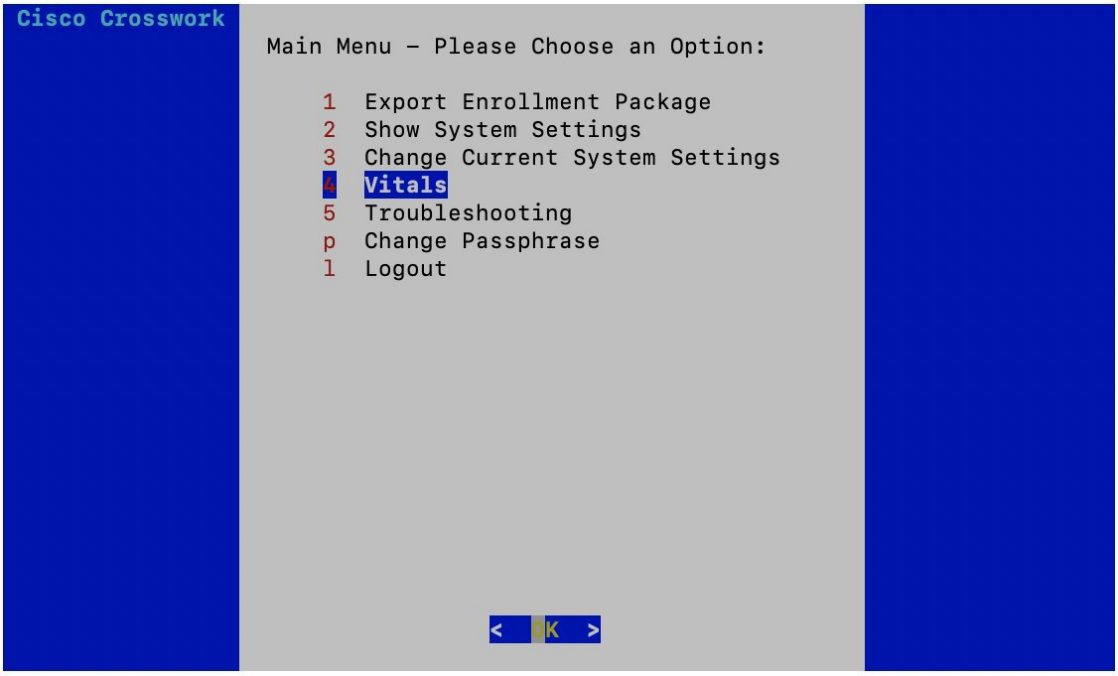

The **Show VM Vitals** menu opens.

Step 2 Select the vital you want to view and click **OK**.

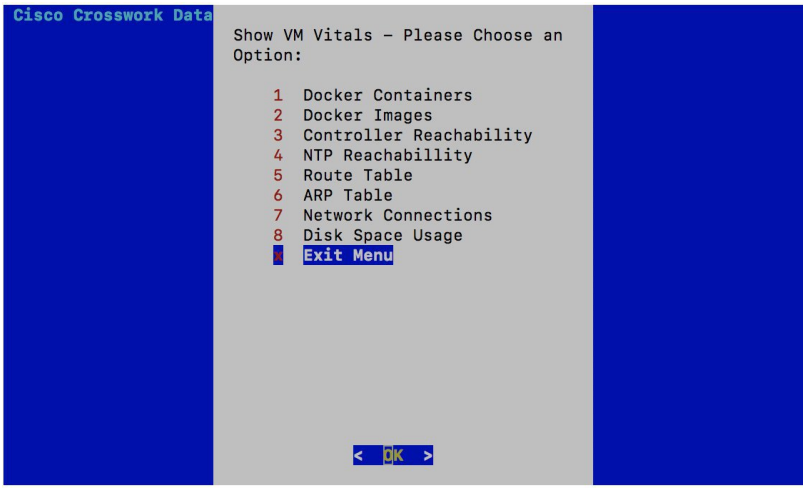

Vital	Description
Docker Containers	Displays the following vitals for the docker containers: • Container ID • Image • Name • Command • Created Time • Status • Port
Docker Images	Displays the following vitals for the docker images: • Repository • Image ID • Created Time • Size • Tag
Controller Reachability	Displays the following vitals for controller reachability: • Default gateway status • Reachability test details (number of packets transmitted and received, packet loss percentage, and time) • DNS server • DNS server status • Reachability test details (number of packets transmitted and received, packet loss percentage, and time) • Controller session status

Vital	Description
NTP Reachability	Displays the following vitals for NTP reachability: • NTP server • Resolved IP Address • Status • Reachability test details (number of packets transmitted and received, packet loss percentage, and time) • Chrony status • Reference ID • System time
Route Table	Displays IPv4 and IPv6 route tables.
ARP Table	Displays ARP tables.
Network Connections	Displays the following vitals for network connections: • Netid • State • Recv-Q • Send-Q • Local Address and Port • Peer Address and Port
Disk Space Usage	Displays the following vitals for disk space usage: • Filesystem • Size • Used space • Available space • Use percentage • Mounted on volume

Cisco Crosswork Data Gateway displays the vitals for the selected item.

After you are done viewing the vitals, press any key to return to the **ShowVM Vitals** menu.

To return to the Main Menu, select **x Exit Menu** and click **OK**.

collector-vitals Service

Cisco Crosswork Data Gateway comprises of various containerized services running on an Ubuntu VM. Its overall health depends on health of each containerized service.

As part of collector vitals, Cisco Crosswork Data Gateway collects host and container metrics and writes them to a container mounted path in vitals.json file and sends it to the Controller.

These vitals of a Cisco Crosswork Data Gateway VM can also be viewed in the Crosswork UI as described in Section: View Cisco Crosswork Data Gateway Instance Health, on page 172.

It collects the following metrics:

Field	Description
Host VM	
Disk Space Used	Percentage of the disk space used for partitions: / /opt/dg/log /var/lib/docker
Disk In/Out	Number of read/write or input/output operations involving a disk for the partitions: / /opt/dg/log /var/lib/docker **Note** This is a cumulative counter, not a delta time series.
CPU Utilization	Amount of actively used CPU and total number of vCPUs.
Load	Load average – is the average system load over a given period of time of 1, 5, and 15 minutes.
Memory	Amount of memory used and available memory.
Network In/Out	The amount of data sent/received in MB for NIC interfaces: eth0 eth1 eth2 **Note** This is a cumulative counter, not a delta time series.
Service Status	

Field	Description
Service	Name of the Cisco Crosswork Data Gateway service.
Status	Status of the service: • Running • Degraded • Error
CPU Utilization	Percentage of actively utilized CPU by the service.
Version	Version of the service deployed.
Memory Used (MB)	Amount of memory being used by the service.
Network In/Out	The amount of data sent/received in MB by the service over its interface. **Note** This is a cumulative counter, not a delta time series.
Disk In/Out	Number of read/write or input/output operations that the service has done involving a disk. **Note** This is a cumulative counter, not a delta time series.

Note

- When either of the following components listed below are not responsive, Cisco Crosswork Data Gateway vitals are not updated:
  - Docker Engine
  - Vitals Monitor
  - Controller Gateway

  The "Collector Vitals" and "Controller Gateway" dockers must be up and running for alerts/vitals to get updated.
- When Vitals Monitor stops, no alerts are added to cdg-alerts.log. This is because the monitor service runs as a part of Vitals Monitors and it doesn't trigger any alerts when Vitals Monitor itself is down.
- Also, the alerts are not added to cdg-alerts.log when Vitals Monitor is running and Controller Gateway is down.

  To access log files, see Run show-tech, on page 232.

Troubleshooting

You can troubleshoot a Cisco Crosswork Data Gateway instance directly from the VM. Cisco Crosswork Data Gateway provides logs of errors, requests to the server, and changes made to the VM and reports any process failures/outages.

To access **Troubleshooting** menu, select **5 Troubleshooting** from the Main Menu and click **OK**, as shown in the following figure:

Cisco Crosswork Data Gateway opens the **Troubleshooting** menu that provides you the following options to troubleshoot your Cisco Crosswork Data Gateway instance:

Note The following figure shows the Troubleshooting Menu corresponding to **dg-admin** user. Few of these options are not available to **dg-oper** user. See Table Table 17: Permissions Per Role, on page 212.

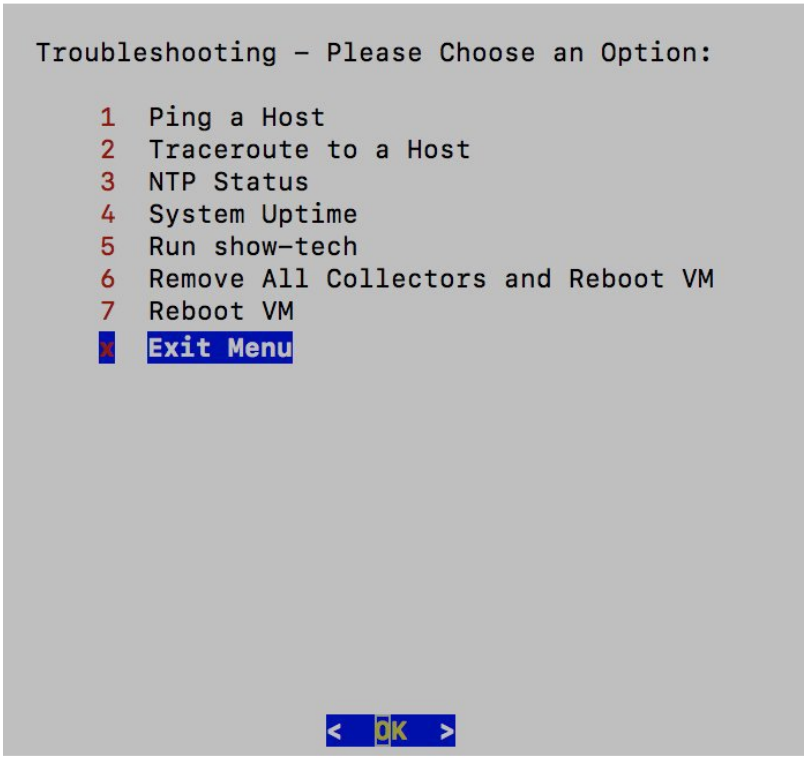

This section contains the following topics:

- Ping a Host, on page 229
- Traceroute to a Host, on page 230
- Check NTP Status, on page 231
- Check System Uptime, on page 231
- Run show-tech, on page 232
- Reboot Crosswork Data Gateway VM, on page 233

Ping a Host

To aid troubleshooting, Cisco Crosswork Data Gateway provides you Ping utility that can be used to check reachability to any IP address.

Step 1 From **Troubleshooting** menu, select **1 Ping a Host** and click **OK**.

Step 2 Enter the ping destination.

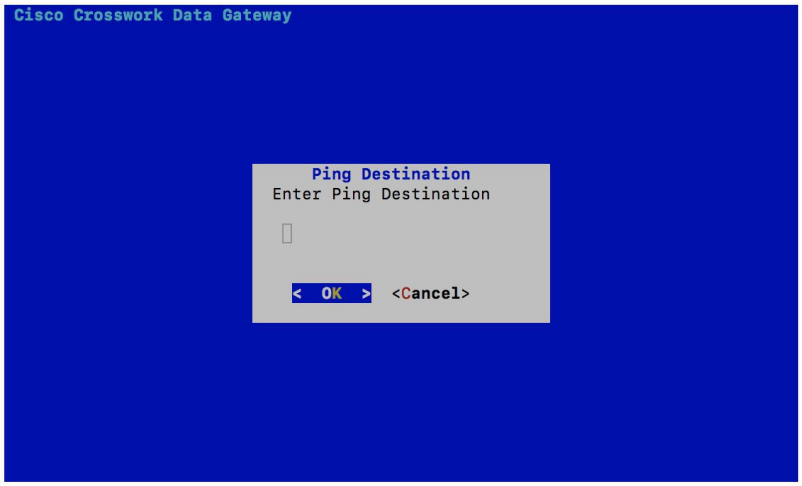

Step 3 Click **OK**.

Cisco Crosswork Data Gateway displays the result of the ping operation.

```
PING 172.23.92.143 (172.23.92.143) 56(84) bytes of data.
64 bytes from 172.23.92.143: icmp_seq=1 ttl=64 time=0.428 ms
64 bytes from 172.23.92.143: icmp_seq=2 ttl=64 time=0.368 ms
64 bytes from 172.23.92.143: icmp_seq=3 ttl=64 time=0.270 ms

64 bytes from 172.23.92.143: icmp_seq=4 ttl=64 time=0.574 ms

64 bytes from 172.23.92.143: icmp_seq=5 ttl=64 time=0.433 ms
64 bytes from 172.23.92.143: icmp_seq=6 ttl=64 time=0.487 ms
^C
--- 172.23.92.143 ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 5107ms
rtt min/avg/max/mdev = 0.270/0.426/0.574/0.097 ms
Press any key to continue
```

Traceroute to a Host

Cisco Crosswork Data Gateway provides **Traceroute to a Host** option to help troubleshoot latency issues. Using this option provides you a rough time estimate for the Cisco Crosswork Data Gateway to reach the controller application.

Step 1 From **Troubleshooting** menu, select **2 Traceroute to a Host** and click **OK**.

Step 2 Enter the traceroute destination.

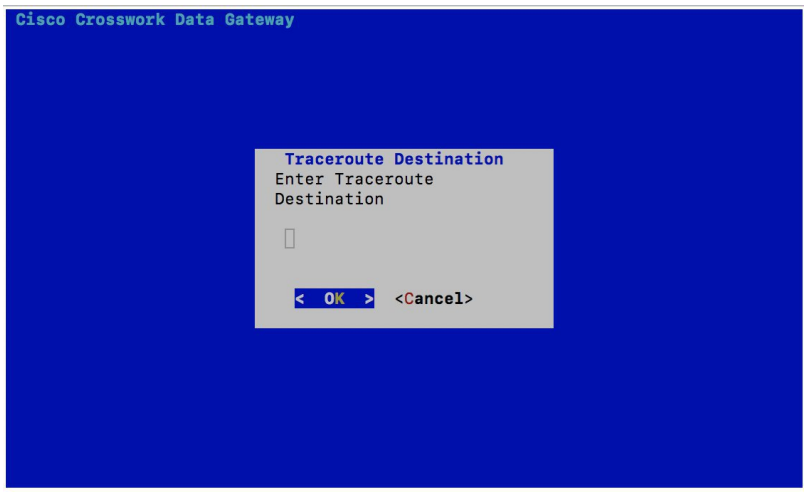

Step 3 Click **OK**.

Check NTP Status

Use this option to check the status of the NTP server.

Step 1 From **Troubleshooting** menu, select **3 NTP Status**.

Step 2 Click **OK**. The Cisco Crosswork Data Gateway displays the NTP server status.

```
Reference ID    : AB442641 (mtv5-ai27-dcm10n-ntp1.cisco.com)
Stratum         : 2
Ref time (UTC)  : Fri Jun 21 04:53:44 2019
System time     : 0.000044881 seconds fast of NTP time
Last offset     : +0.000057586 seconds
RMS offset      : 0.000080841 seconds
Frequency       : 21.559 ppm slow
Residual freq   : +0.009 ppm
Skew            : 0.144 ppm
Root delay      : 0.002095408 seconds
Root dispersion : 0.001190380 seconds
Update interval : 2062.6 seconds
Leap status     : Normal
Press any key to continue
```

Check System Uptime

Use this option to check system uptime.

Step 1 From **Troubleshooting** menu, select **4 System Uptime**.

Step 2 Click **OK**. The Crosswork Data Gateway displays the system uptime.

```
 05:11:55 up 3 days,  1:49,  1 user,  load average: 0.18, 0.12, 0.10
Press any key to continue
```

Run show-tech

Cisco Crosswork Data Gateway provides the option **show_tech** to export its log files to a user-defined SCP destination.

The collected data includes the following:

- Logs of all the Data Gateway components running on docker containers
- VM Vitals

It creates a tarball in the directory where it is executed. The output is a tarball named `CDG-<CDG-version>-year-month-day--hour-minute-second-*.tar.bz2`

The execution of this command may take several minutes depending on the state of Crosswork Data Gateway.

Step 1 From **Troubleshooting** menu, select **5 Show-tech** and click **OK**.

Step 2 Enter the destination to save the tarball containing logs and vitals.

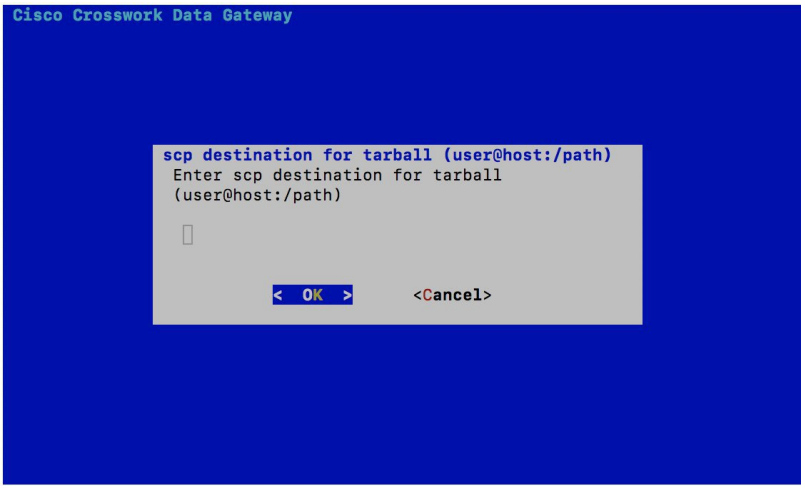

Step 3 Enter your SCP passphrase and click **OK**.

Reboot Crosswork Data Gateway VM

Note This task can only be performed by **dg-admin** user.

Crosswork Data Gateway gives you two options to reboot the VM:

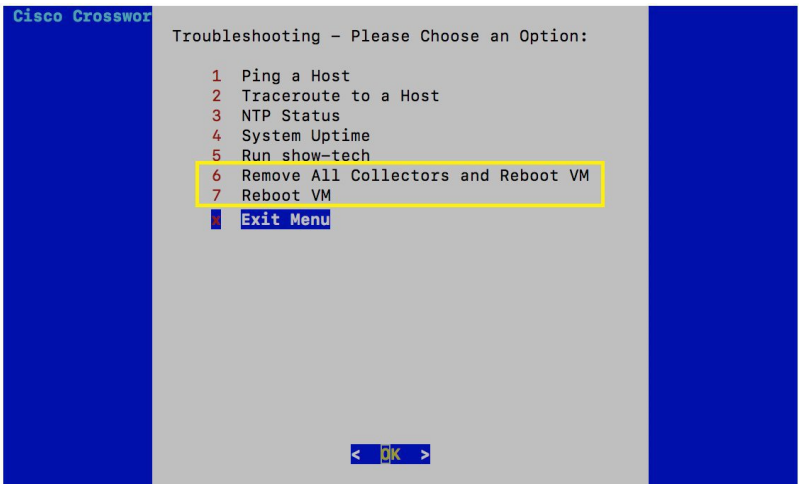

- **Remove All Collectors and Reboot VM**: Select this option from the **Troubleshooting** menu if you want to remove all the collectors (functional images) and reboot VM.
- **Reboot VM**: Select this option from the **Troubleshooting** menu for a normal reboot.